

Efficient Methods for Non-stationary Online Learning

Peng Zhao

Yan-Feng Xie

Lijun Zhang

Zhi-Hua Zhou

ZHAOP@LAMDA.NJU.EDU.CN

XIEYF@LAMDA.NJU.EDU.CN

ZHANGLJ@LAMDA.NJU.EDU.CN

ZHOUGH@LAMDA.NJU.EDU.CN

National Key Laboratory for Novel Software Technology, Nanjing University, China

School of Artificial Intelligence, Nanjing University, China

Editor: Alekh Agarwal

Abstract

Non-stationary online learning has drawn much attention in recent years. In particular, *dynamic regret* and *adaptive regret* are proposed as two principled performance measures for online convex optimization in non-stationary environments. To optimize them, a two-layer online ensemble is usually deployed due to the inherent uncertainty of non-stationarity, in which multiple base-learners are maintained and a meta-algorithm is employed to track the best one on the fly. However, the two-layer structure raises concerns about computational complexity — such methods typically maintain $\mathcal{O}(\log T)$ base-learners simultaneously for a T -round online game and thus perform multiple projections onto the feasible domain per round, which becomes the computational bottleneck when the domain is complicated. In this paper, we present efficient methods for optimizing dynamic regret and adaptive regret that reduce the number of projections per round from $\mathcal{O}(\log T)$ to 1. The proposed algorithms require only one gradient query and one function evaluation at each round. Our technique hinges on the reduction mechanism developed in parameter-free online learning and requires non-trivial modifications for non-stationary online methods. Furthermore, we study an even stronger measure, namely “interval dynamic regret”, and reduce the number of projections per round from $\mathcal{O}(\log^2 T)$ to 1 for minimizing it. Our reduction demonstrates broad generality and applies to two important applications: online stochastic control and online principal component analysis, resulting in methods that are both efficient and optimal. Finally, empirical studies verify our theoretical findings.

1. Introduction

Classic online learning focuses on minimizing the static regret, which evaluates the online learner’s performance against the best fixed decision in hindsight (Hazan, 2016). However, in many real-world applications, the environments are often non-stationary. In such scenarios, minimizing static regret becomes less attractive, since it would be unrealistic to assume the existence of a single decision behaved satisfactorily throughout the entire time horizon.

To address the limitation, in recent years, researchers have studied more strengthened performance measures to facilitate online algorithms with the capability of handling non-stationarity. In particular, dynamic regret (Zinkevich, 2003; Zhang et al., 2018a) and adaptive regret (Hazan and Seshadhri, 2009; Daniely et al., 2015) are proposed as two principled metrics to guide the algorithm design. We focus on the online convex optimization

(OCO) setting (Hazan, 2016), which can be deemed as a game between the learner and the environments. At each round $t \in [T]$, the learner submits her decision $\mathbf{x}_t \in \mathcal{X}$ from a convex feasible domain $\mathcal{X} \subseteq \mathbb{R}^d$ and simultaneously environments choose a convex function $f_t : \mathcal{X} \mapsto \mathbb{R}$, and subsequently the learner suffers an instantaneous loss $f_t(\mathbf{x}_t)$.

1.1 Dynamic Regret and Adaptive Regret

Dynamic regret is proposed by Zinkevich (2003) to compare the online learner’s performance against a sequence of *any* feasible comparators $\mathbf{u}_1, \dots, \mathbf{u}_T \in \mathcal{X}$. Formally, it is defined as

$$\text{D-REG}_T(\mathbf{u}_1, \dots, \mathbf{u}_T) = \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}_t). \quad (1)$$

Dynamic regret minimization enables the learner to track changing comparators. A favorable dynamic regret bound should scale with a certain non-stationarity measure dependent on the comparators such as the path length $P_T = \sum_{t=2}^T \|\mathbf{u}_t - \mathbf{u}_{t-1}\|_2$. Notably, the classic static regret can be treated as a special case of dynamic regret by specifying the comparators as the best fixed decision in hindsight.

Adaptive regret is proposed by Hazan and Seshadhri (2009) and further strengthened by Daniely et al. (2015), which measures the regret over *any* interval $I = [r, s] \subseteq [T]$ with a length of $\tau = |I|$ and hence is also referred to as the *interval regret*. The specific definition is

$$\text{A-REG}_T(|I|) = \max_{[r, r+\tau-1] \subseteq [T]} \left\{ \sum_{t=r}^{r+\tau-1} f_t(\mathbf{x}_t) - \min_{\mathbf{u} \in \mathcal{X}} \sum_{t=r}^{r+\tau-1} f_t(\mathbf{u}) \right\}. \quad (2)$$

Since the minimizers of different intervals can be different, adaptive regret minimization also ensures the capability of competing with changing comparators. A desired adaptive regret bound should be as close as the minimax static regret of this interval. Algorithms with adaptive regret matching static regret of this interval up to logarithmic terms in T are referred to as strongly adaptive (Daniely et al., 2015). Moreover, it can be observed that adaptive regret can include the static regret when choosing $I = [T]$.

Substantial research has been devoted to optimizing these two measures, including algorithms for dynamic regret (Zinkevich, 2003; Hall and Willett, 2013; György and Szepesvári, 2016; Zhang et al., 2018a; Zhao et al., 2020, 2021, 2024; Baby and Wang, 2021; Jacobsen and Cutkosky, 2022; Zhao et al., 2022a; Bai et al., 2022; Zhang et al., 2023; Li et al., 2024; Zhang et al., 2025) and adaptive regret (Hazan and Seshadhri, 2009; Daniely et al., 2015; Adamskiy et al., 2016; Jun et al., 2017; Zhang et al., 2018b, 2019; Cutkosky, 2020; Zhang et al., 2021; Yang et al., 2024). It is worth noting that the relationship between dynamic regret and adaptive regret in online convex optimization (OCO) remains largely unclear (Zhang, 2020, Section 5), even though a black-box reduction from dynamic regret to adaptive regret has been established for the simpler prediction-with-expert-advice setting (i.e., online linear optimization over simplex) (Luo and Schapire, 2015, Theorem 4). As a result, the two measures are typically developed independently.

To achieve the best of both worlds, several studies (Zhang et al., 2020; Cutkosky, 2020) have proposed optimizing both measures simultaneously by optimizing an even strengthened

metric called *interval dynamic regret*, defined as

$$\text{INTERVAL-D-REG}_T(|I|; \mathbf{u}_1, \dots, \mathbf{u}_T) = \max_{[r, r+\tau-1] \subseteq [T]} \left\{ \sum_{t=r}^{r+\tau-1} f_t(\mathbf{x}_t) - \sum_{t=r}^{r+\tau-1} f_t(\mathbf{u}_t) \right\}. \quad (3)$$

This measure requires the online algorithm to compete with a sequence of time-varying comparators over any interval.

1.2 Two-layer Online Ensemble and Projection Complexity Issue

The fundamental challenge of optimizing the above non-stationary regret measures is the uncertainty of environmental non-stationarity. Concretely, to ensure robustness in unknown environments, dynamic regret aims to compete with *any* feasible comparator sequence, while adaptive regret examines the local performance over *any* intervals; with interval dynamic regret is even more ambitious. The unknown comparators or/and unknown intervals bring considerable uncertainty to online optimization. To address the issue, a two-layer online ensemble structure is usually deployed to optimize the measures (Zhao et al., 2024), where a set of base-learners are maintained to handle the different possibilities of online environments and a meta-algorithm is employed to combine them all and track the unknown best one. Such a meta-base framework successfully achieves many state-of-the-art results, including the $\mathcal{O}(\sqrt{T(1+P_T)})$ dynamic regret (Zhang et al., 2018a) and the $\mathcal{O}(\sqrt{(F_T^{\mathbf{u}} + P_T)(1+P_T)})$ small-loss dynamic regret for smooth functions (Zhao et al., 2020), where $P_T = \sum_{t=2}^T \|\mathbf{u}_t - \mathbf{u}_{t-1}\|_2$ is the path length and $F_T^{\mathbf{u}} = \sum_{t=1}^T f_t(\mathbf{u}_t)$ is the cumulative loss of comparators; as well as the $\mathcal{O}(\sqrt{|I| \log T})$ adaptive regret (Jun et al., 2017) and the $\mathcal{O}(\sqrt{F_I \log F_I \log F_T})$ small-loss adaptive regret for smooth functions (Zhang et al., 2019) for any interval $I = [r, s] \subseteq [T]$, where $F_I = \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=r}^s f_t(\mathbf{x})$ and $F_T = \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T f_t(\mathbf{x})$. Besides, an $\mathcal{O}(\sqrt{|I|(\log T + P_I)})$ interval dynamic regret is also achieved by a two-layer (or even three-layer) structure (Zhang et al., 2020), where $P_I = \sum_{t=r}^s \|\mathbf{u}_t - \mathbf{u}_{t-1}\|_2$ is the path length over the interval.

The two-layer online ensemble methods have demonstrated great effectiveness in tackling non-stationary online environments, whereas the gain is at the price of heavier computations than the methods for minimizing static regret. While it is believed that additional computations are necessary for more robustness, we are wondering whether it is possible to pay for a “minimal” computation overhead for adapting to the non-stationarity. To this end, we focus on the popular first-order online methods and aim to streamline unnecessary computations while retaining the same regret guarantees. Arguably, the most computationally expensive step of each round is the projection onto the convex feasible domain, namely, the projection operation $\Pi_{\mathcal{X}}[\mathbf{y}] = \arg \min_{\mathbf{x} \in \mathcal{X}} \|\mathbf{x} - \mathbf{y}\|_2$ for a convex set $\mathcal{X} \subseteq \mathbb{R}^d$. Typical two-layer non-stationary online algorithms require maintaining $N = \mathcal{O}(\log T)$ base-learners simultaneously to cover the possibility of unknown environments. Define the *projection complexity* of online methods as the number of projections onto the feasible domain per round. Then, those non-stationary methods suffer an $\mathcal{O}(\log T)$ projection complexity, whereas standard online methods for static regret minimization require only 1 projection onto the feasible domain per round such as online gradient descent (Zinkevich, 2003).

1.3 Our Contributions and Techniques

In this paper, we design a generic mechanism to reduce the projection complexity of many existing non-stationary methods from $\mathcal{O}(\log T)$ to 1 *without sacrificing regret optimality*, thereby matching the projection complexity of stationary methods. Our reduction is inspired by recent advances in parameter-free online learning (Cutkosky and Orabona, 2018; Mhammedi et al., 2019). The idea is simple: we reduce the original problem learned in the feasible domain \mathcal{X} to an alternative one learned in a *surrogate domain* $\mathcal{Y} \supseteq \mathcal{X}$ where projection is much cheaper, e.g., simply choosing \mathcal{Y} as a properly scaled Euclidean ball; and moreover, a carefully designed *surrogate loss* is necessary for the alternative problem to preserve the regret optimality. We reveal that a necessary condition for our reduction mechanism to deploy and reduce the projection complexity is that the non-stationary online algorithm shall *query the gradient once and evaluate the function value once per round*. Several existing algorithms for obtaining the worst-case (problem-independent) dynamic regret or adaptive regret already satisfy the requirements, and our reduction immediately yields efficient counterparts with the same regret guarantees but only one projection per round. However, many non-stationary algorithms, particularly those designed for small-loss bounds, do not satisfy this condition, requiring non-trivial efforts to make them compatible. To this end, we develop a family of algorithms that achieve worst-case/small-loss dynamic regret and adaptive regret with with one projection, one gradient query, and one function evaluation per round.

Despite that the reduction mechanism of this paper has been studied in parameter-free online learning, applying it to non-stationary online learning requires new ideas and non-trivial modifications. Here we highlight the technical innovation. The main challenge comes from the reduction condition mentioned earlier: since the surrogate loss involves the projection operation, our reduction requires the algorithm query one gradient and evaluate one function value at each round. However, many non-stationary algorithms do not satisfy the requirement, which is to be contrasted to the parameter-free algorithms such as MetaGrad (van Erven and Koolen, 2016) or its variants (Wang et al., 2019; Mhammedi et al., 2019) that naturally satisfy the condition. For example, the SACS algorithm (Zhang et al., 2019) enjoys the best known small-loss adaptive regret, yet the method requires N gradient queries and $N + 1$ function evaluations at each round, where $N = \mathcal{O}(\log T)$ is the number of base-learners. Thus, we have to dig into the algorithm and modify it to fit our reduction. First, we replace their meta-algorithm with Adapt-ML-Prod (Gaillard et al., 2014), an expert-tracking algorithm with a *second-order* regret with excess losses to accommodate the linearized loss that is used to ensure one gradient query per round. Second, we introduce a sequence of *time-varying* thresholds to adaptively determine the problem-dependent geometric covers in contrast to a fixed threshold used in their method. In particular, we register the cumulative loss of the final decisions rather than the base-learner’s one to compare it with the changing thresholds, which renders the design of one function value evaluation per round and also turns out to be crucial for achieving an improved small-loss bound that can recover the best known worst-case adaptive adaptive regret (by contrast, SACS cannot obtain optimal worst-case adaptive regret). Our algorithm only requires one projection/gradient query/function evaluation at each round, substantially improving the efficiency of SACS algorithm that requires N projections/gradient queries/function evaluations per round.

Table 1: Summary of different *small-loss* regret results for smooth functions. The third column lists the number of projections required per iteration. The fourth and fifth columns indicate the number of gradient and function value queries per round, respectively, with “–” denoting that function values are not required. The sixth column presents the regret bounds for different algorithms, where $F_I = \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=r}^s f_t(\mathbf{x})$, $F_I^{\mathbf{u}} = \sum_{t=r}^s f_t(\mathbf{u}_t)$, and $P_I = \sum_{t=r}^s \|\mathbf{u}_t - \mathbf{u}_{t-1}\|_2$ are problem-dependent quantities over the interval $I \subseteq [1, T]$. For simplicity, we use F_T to denote $F_{[T]}$ (similarly defined for P_T and $F_T^{\mathbf{u}}$).

Measure	Reference	# Proj	# Grad	# Value	Regret Bound
D-Reg	Zhao et al. (2024)	$\mathcal{O}(\log T)$	1	–	$\mathcal{O}(\sqrt{(F_T^{\mathbf{u}} + P_T)(1 + P_T)})$
	Theorem 3 (ours)	1	1	–	$\mathcal{O}(\sqrt{(F_T^{\mathbf{u}} + P_T)(1 + P_T)})$
A-Reg	Zhang et al. (2019)	$\mathcal{O}(\log T)$	$\mathcal{O}(\log T)$	$\mathcal{O}(\log T)$	$\mathcal{O}(\sqrt{F_I \log F_I \log F_T})$
	Theorem 4 (ours)	1	1	1	$\mathcal{O}(\min\{\sqrt{F_I \log F_I \log F_T}, \sqrt{ I \log T}\})$
Interval D-Reg	Zhang et al. (2020)	$\mathcal{O}(\log^2 T)$	$\mathcal{O}(\log^2 T)$	$\mathcal{O}(\log^2 T)$	$\mathcal{O}(\sqrt{ I (\log T + P_I)})$
	Theorem 5 (ours)	1	1	1	$\mathcal{O}(\sqrt{(F_I^{\mathbf{u}} + P_I)(\log(F_I^{\mathbf{u}} + P_I) \cdot \log(F_T^{\mathbf{u}} + P_T) + P_I)})$

In this paper, extending our preliminary conference version (Zhao et al., 2022c), we further study a stronger measure, *interval dynamic regret*, and develop an efficient method that reduces its projection complexity from $\mathcal{O}(\log^2 T)$ to 1 using our reduction mechanism. Analyzing interval dynamic regret requires techniques similar to those for adaptive regret, along with additional refinements in the geometric covering intervals and base-algorithm to handle the more complex structure. Table 1 summarizes our main results in comparison with existing ones, in which we focus on the small-loss bounds for convex and smooth functions. Our reduction mechanism demonstrates great generality and can be applied to two important applications: (i) online non-stochastic control (Agarwal et al., 2019a), and (ii) online principal component analysis (Warmuth and Kuzmin, 2008). In both applications, the projection operation is complicated and time-consuming, which becomes even worse when leveraging these methods to handle non-stationary environments, since typically $\mathcal{O}(\log T)$ projections per round are required by following previous methods (Zhang et al., 2019; Zhao et al., 2023). We explore the structure of these two problems and propose efficient methods with 1 projection per round to solve them under non-stationary environments, while ensuring optimal theoretical guarantees.

1.4 Assumptions

In this part, we list several standard assumptions used in OCO (Shalev-Shwartz, 2012; Hazan, 2016). Notably, not all these assumptions are always required. We will explicitly state the requirements in the theorem.

Assumption 1 (bounded gradient). The norm of the gradients of online functions over the domain \mathcal{X} is bounded by G , i.e., $\|\nabla f_t(\mathbf{x})\|_2 \leq G$, for all $\mathbf{x} \in \mathcal{X}$ and $t \in [T]$.

Assumption 2 (bounded domain). The domain $\mathcal{X} \subseteq \mathbb{R}^d$ contains the origin $\mathbf{0}$, and the diameter of the domain \mathcal{X} is at most D , i.e., $\|\mathbf{x} - \mathbf{x}'\|_2 \leq D$ for any $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$.

Assumption 3 (non-negativity and smoothness). All the online functions are non-negative and L -smooth, i.e., for $t \in [T]$, the online function $f_t : \mathbb{R}^d \mapsto \mathbb{R}$ satisfies $\|\nabla f_t(\mathbf{x}) - \nabla f_t(\mathbf{x}')\|_2 \leq L\|\mathbf{x} - \mathbf{x}'\|_2$ for any $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^d$.

Note that in Assumption 3 we assume that the online functions are smooth over the *entire space* \mathbb{R}^d . This is primarily for technical reasons, as it ensures the self-bounding property (described in Lemma 20) holds when deriving small-loss regret bounds (Srebro et al., 2010; Orabona et al., 2012; Zhang et al., 2019; Zhao et al., 2020). The constraint on the feasible domain remains valid, as it can be understood by noting that, while the online functions may be defined over \mathbb{R}^d , the decisions are only feasible within a subset $\mathcal{X} \subseteq \mathbb{R}^d$, which is also the case in our two applications as discussed in Section 5.

In fact, it is possible to relax Assumption 3 by requiring it to hold only over an extended domain \mathcal{X}' , which is slightly larger than the original domain \mathcal{X} . Specifically, this extended domain is defined as $\mathcal{X}' = \{\mathbf{x} + \mathbf{y} \mid \mathbf{x} \in \mathcal{X}, \mathbf{y} \in \mathbb{R}^d, \|\mathbf{y}\|_2 \leq G/L\}$; in other words, \mathcal{X}' is the Minkowski sum of the original domain \mathcal{X} and a small Euclidean ball.

1.5 Paper Outline

The remainder of the paper is structured as follows. In Section 2, we delineate the reduction mechanism and illustrate its application to dynamic regret minimization. Section 3 offers efficient methods for optimizing adaptive regret. Subsequently, Section 4 gives the results for optimizing an even stronger performance measure – interval dynamic regret. In Section 5 we provide two applications of our proposed reduction mechanism for non-stationary online learning. Experimental validations are reported in Section 6. Finally, we conclude the paper and make several discussions in Section 7. All the proofs and omitted details for algorithms are deferred to the appendices.

2. The Reduction Mechanism and Dynamic Regret Minimization

We start from the dynamic regret minimization. First, we briefly review existing methods in Section 2.1, and then present our reduction mechanism and illustrate how to apply it to reducing the projection complexity of dynamic regret methods in Section 2.2.

2.1 A Brief Review of Dynamic Regret Minimization

Zhang et al. (2018a) propose a two-layer online algorithm called Ader with an $\mathcal{O}(\sqrt{T(1 + P_T)})$ dynamic regret, which is proved to be minimax optimal for convex functions. Ader maintains a group of base-learners, each performing online gradient descent (OGD) (Zinkevich, 2003) with a customized step size specified by the pool $\mathcal{H} = \{\eta_1, \dots, \eta_N\}$, and then uses a meta-algorithm to combine them all. Denoted by $\mathcal{B}_1, \dots, \mathcal{B}_N$ the N base-learners. For each $i \in [N]$, the base-learner \mathcal{B}_i updates by

$$\mathbf{x}_{t+1,i} = \Pi_{\mathcal{X}}[\mathbf{x}_{t,i} - \eta_i \nabla f_t(\mathbf{x}_t)], \quad (4)$$

where $\eta_i \in \mathcal{H}$ is the associated step size and $\Pi_{\mathcal{X}}[\cdot]$ denotes the projection onto the feasible domain \mathcal{X} with $\Pi_{\mathcal{X}}[\mathbf{y}] = \arg \min_{\mathbf{x} \in \mathcal{X}} \|\mathbf{y} - \mathbf{x}\|_2$. Notably, all the base-learners share the same gradient $\nabla f_t(\mathbf{x}_t)$ rather than using their individual one $\nabla f_t(\mathbf{x}_{t,i})$. This is because

Ader optimizes the linearized loss $\ell_t(\mathbf{x}) = \langle \nabla f_t(\mathbf{x}_t), \mathbf{x} \rangle$, which enjoys the benign property of $\nabla \ell_t(\mathbf{x}_{t,i}) = \nabla f_t(\mathbf{x}_t)$ for all $i \in [N]$.¹

Furthermore, the meta-algorithm evaluates each base-learner by the linearized loss $\ell_t(\mathbf{x}_{t,i}) = \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_{t,i} \rangle$ and updates the weight vector $\mathbf{p}_{t+1} \in \Delta_N$ by the Hedge algorithm (Freund and Schapire, 1997), namely,

$$p_{t+1,i} = \frac{p_{t,i} \exp(-\varepsilon \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_{t,i} \rangle)}{\sum_{j=1}^N p_{t,j} \exp(-\varepsilon \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_{t,j} \rangle)}, \quad \forall i \in [N], \quad (5)$$

where $\varepsilon > 0$ is meta-algorithm's learning rate. The final prediction is obtained by $\mathbf{x}_{t+1} = \sum_{i=1}^N p_{t+1,i} \mathbf{x}_{t+1,i}$. The learner submits \mathbf{x}_{t+1} and receives the loss $f_{t+1}(\mathbf{x}_{t+1})$ and the gradient $\nabla f_{t+1}(\mathbf{x}_{t+1})$ as the feedback of this round. With a suitable step size pool \mathcal{H} and learning rate $\varepsilon = \Theta(\sqrt{(\ln N)/T})$ with $N = \mathcal{O}(\log T)$, Ader enjoys an $\mathcal{O}(\sqrt{T(1+P_T)})$ dynamic regret using 1 gradient query per round (Zhang et al., 2018a, Theorem 4).

For smooth functions, Zhao et al. (2024) demonstrate that an $\mathcal{O}(\sqrt{(F_T^{\mathbf{u}} + P_T)(1+P_T)})$ small-loss dynamic regret can be achieved where $F_T^{\mathbf{u}} = \sum_{t=1}^T f_t(\mathbf{u}_t)$ is the cumulative loss of the dynamic comparators. This bound safeguards the minimax rate in the worst case, while it can be much smaller than $\mathcal{O}(\sqrt{T(1+P_T)})$ bound in the benign environments.

2.2 The Reduction Mechanism for Reducing Projection Complexity

As demonstrated in the update (4), all the base-learners require projecting the intermediate solution onto the domain \mathcal{X} to ensure feasibility. As a result, $\mathcal{O}(\log T)$ projections are required at each round, which is generally time-consuming particularly when the domain \mathcal{X} is complicated. To address so, we present a generic reduction mechanism for reducing the projection complexity and apply it to dynamic regret methods. Our reduction builds upon the seminal work (Cutkosky and Orabona, 2018) and a further refined result (Cutkosky, 2020), who propose a black-box reduction from constrained online learning to the unconstrained setting (or another constrained problem with a larger domain).

Reduction mechanism. Given an algorithm for non-stationary online learning **Algo** whose projection complexity is $\mathcal{O}(\log T)$, our reduction mechanism builds on it to yield an algorithm **Efficient-Algo** with 1 projection onto \mathcal{X} per round and retaining the same order of regret. The central idea is to replace expensive projections onto the original domain \mathcal{X} with other much cheaper projections. To this end, we introduce a *surrogate domain* \mathcal{Y} defined as the minimum Euclidean ball containing the feasible domain \mathcal{X} , i.e., $\mathcal{Y} = \{\mathbf{x} \mid \|\mathbf{x}\|_2 \leq D\} \supseteq \mathcal{X}$. Then, the reduced algorithm **Algo** works on \mathcal{Y} whose projection can be realized by a simple rescaling. More importantly, to avoid regret degeneration, it is necessary to carefully construct the surrogate loss $g_t : \mathcal{Y} \mapsto \mathbb{R}$ as

$$g_t(\mathbf{y}) = \langle \nabla f_t(\mathbf{x}_t), \mathbf{y} \rangle - \mathbb{1}_{\{\langle \nabla f_t(\mathbf{x}_t), \mathbf{v}_t \rangle < 0\}} \cdot \langle \nabla f_t(\mathbf{x}_t), \mathbf{v}_t \rangle \cdot S_{\mathcal{X}}(\mathbf{y}), \quad (6)$$

where $S_{\mathcal{X}}(\mathbf{y}) = \inf_{\mathbf{x} \in \mathcal{X}} \|\mathbf{y} - \mathbf{x}\|_2$ is the distance function to \mathcal{X} and $\mathbf{v}_t = (\mathbf{y}_t - \mathbf{x}_t) / \|\mathbf{y}_t - \mathbf{x}_t\|_2$ is the vector indicating the projection direction.

The main protocol of our reduction is presented as follows. The input includes original functions $\{f_t\}_{t=1}^T$, the feasible domain \mathcal{X} , and the reduced algorithm **Algo**.

1. We refer to the improved approach presented in Section 3.4 of Zhang et al. (2018a).

- 1: **for** $t = 1, \dots, T$ **do**
- 2: receive the gradient information $\nabla f_t(\mathbf{x}_t)$;
- 3: construct the surrogate loss $g_t : \mathcal{Y} \mapsto \mathbb{R}$ according to [Eq. \(6\)](#);
- 4: obtain the intermediate prediction $\mathbf{y}_{t+1} = \text{Algo}(g_t(\cdot), \mathbf{y}_t, \mathcal{Y})$;
- 5: submit the final prediction $\mathbf{x}_{t+1} = \Pi_{\mathcal{X}}[\mathbf{y}_{t+1}]$;
- 6: **end for**

Our reduction enjoys the regret safeness due to the benign properties of surrogate loss.

Theorem 1 (Theorem 2 of [Cutkosky \(2020\)](#)). *The surrogate loss $g_t : \mathcal{Y} \mapsto \mathbb{R}$ defined in [\(6\)](#) is convex. Moreover, we have $\|\nabla g_t(\mathbf{y}_t)\|_2 \leq \|\nabla f_t(\mathbf{x}_t)\|_2$ and for any $\mathbf{u}_t \in \mathcal{X}$*

$$\langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u}_t \rangle \leq g_t(\mathbf{y}_t) - g_t(\mathbf{u}_t) \leq \langle \nabla g_t(\mathbf{y}_t), \mathbf{y}_t - \mathbf{u}_t \rangle. \quad (7)$$

The theorem shows the convexity of the surrogate loss $g_t(\mathbf{y})$ and we thus have $f_t(\mathbf{x}_t) - f_t(\mathbf{u}_t) \leq \langle \nabla g_t(\mathbf{y}_t), \mathbf{y}_t - \mathbf{u}_t \rangle$, which implies that it suffices to optimize the linearized loss $\ell_t(\mathbf{y}) = \langle \nabla g_t(\mathbf{y}_t), \mathbf{y} \rangle$. Moreover, the property $\|\nabla g_t(\mathbf{y}_t)\|_2 \leq \|\nabla f_t(\mathbf{x}_t)\|_2$ established in [Theorem 1](#), is crucial for achieving the small-loss bound, as will be demonstrated later in [Theorem 3](#). Furthermore, we have the following lemma that specifies the gradient calculation.

Lemma 1. *For any $\mathbf{y} \in \mathcal{Y}$, $\nabla g_t(\mathbf{y}) = \nabla f_t(\mathbf{x}_t)$ when $\langle \nabla f_t(\mathbf{x}_t), \mathbf{v}_t \rangle \geq 0$; and $\nabla g_t(\mathbf{y}) = \nabla f_t(\mathbf{x}_t) - \langle \nabla f_t(\mathbf{x}_t), \mathbf{v}_t \rangle \cdot (\mathbf{y} - \Pi_{\mathcal{X}}[\mathbf{y}]) / \|\mathbf{y} - \Pi_{\mathcal{X}}[\mathbf{y}]\|_2$ when $\langle \nabla f_t(\mathbf{x}_t), \mathbf{v}_t \rangle < 0$. Here $\mathbf{v}_t = (\mathbf{y}_t - \mathbf{x}_t) / \|\mathbf{y}_t - \mathbf{x}_t\|_2$. In particular, $\nabla g_t(\mathbf{y}_t) = \nabla f_t(\mathbf{x}_t) - \langle \nabla f_t(\mathbf{x}_t), \mathbf{v}_t \rangle \cdot \mathbf{v}_t$ when $\langle \nabla f_t(\mathbf{x}_t), \mathbf{v}_t \rangle < 0$.*

Reduction requirements. An important necessary condition for the reduction is to require the reduced algorithm satisfying *one gradient query* and *one function evaluation* at each round. Indeed, the reduction essentially updates according to the surrogate loss $\{g_t\}_{t=1}^T$. Note that the definition of surrogate loss involves the distance function $S_{\mathcal{X}}(\mathbf{y})$, see [Eq. \(6\)](#). Thus, each evaluation of $g_t(\mathbf{y})$ leads to one projection onto \mathcal{X} due to the calculation of $S_{\mathcal{X}}(\mathbf{y})$. Similarly, each gradient query of $\nabla g_t(\mathbf{y})$ also contributes to one projection, see [Lemma 1](#) for details. To summarize, we can use the reduction to ensure a 1 projection complexity, only when the reduced algorithm satisfies the requirements of one gradient query and one function evaluation per round. Below, we demonstrate the usage of our reduction mechanism for two methods of dynamic regret minimization that satisfy the conditions, including the worst-case method ([Zhang et al., 2018a](#)) and the small-loss method ([Zhao et al., 2024](#)).

Application to dynamic regret minimization. [Algorithm 1](#) summarizes the main procedures of our efficient methods for optimizing dynamic regret, which is an instance of the reduction mechanism by picking [Algo](#) as [Ader](#) ([Zhang et al., 2018a](#)). More specifically, [Lines 6 – 8](#) are essentially performing [Ader](#) algorithm using the surrogate loss $\{g_t\}_{t=1}^T$ over the surrogate domain \mathcal{Y} . Note that the base update in [Line 6](#) is essentially performing OGD with projection onto \mathcal{Y} , a scaled Euclidean ball, and thus the projection admits a simple closed form. The overall algorithm requires projecting onto \mathcal{X} only once per round, see [Line 9](#). Our method provably retains the same dynamic regret.

Theorem 2. *Set the step size pool as $\mathcal{H} = \{\eta_i = 2^{i-1}(D/G)\sqrt{5/(2T)} \mid i \in [N]\}$ with $N = \lceil 2^{-1} \log_2(1 + 2T/5) \rceil + 1$ and the learning rate as $\varepsilon = \sqrt{(\ln N)/(1 + G^2 D^2 T)}$. Under [Assumptions 1](#) and [2](#), our algorithm requires one projection onto \mathcal{X} per round and enjoys*

$$\sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}_t) \leq \mathcal{O}\left(\sqrt{T(1 + P_T)}\right). \quad (8)$$

Algorithm 1 Efficient Algorithm for Minimizing Dynamic Regret

Input: step size pool $\mathcal{H} = \{\eta_1, \dots, \eta_N\}$, learning rate of meta-algorithm ε_t (or simply a fixed one $\varepsilon_t = \varepsilon$).

- 1: Initialization: let \mathbf{x}_1 and $\{\mathbf{y}_{1,i}\}_{i=1}^N$ be any point in \mathcal{X} ; $\forall i \in [N], p_{1,i} = 1/N$.
- 2: **for** $t = 1$ **to** T **do**
- 3: Receive the gradient information $\nabla f_t(\mathbf{x}_t)$.
- 4: Construct the surrogate loss $g_t : \mathcal{Y} \mapsto \mathbb{R}$ according to Eq. (6).
- 5: Compute the gradient $\nabla g_t(\mathbf{y}_t)$ according to Lemma 1.
- 6: For each $i \in [N]$, the base-learner \mathcal{B}_i produces the local decision by

$$\hat{\mathbf{y}}_{t+1,i} = \mathbf{y}_{t,i} - \eta_i \nabla g_t(\mathbf{y}_t), \quad \mathbf{y}_{t+1,i} = \hat{\mathbf{y}}_{t+1,i} \left(\mathbb{1}_{\{\|\hat{\mathbf{y}}_{t+1,i}\|_2 \leq D\}} + \frac{D}{\|\hat{\mathbf{y}}_{t+1,i}\|_2} \cdot \mathbb{1}_{\{\|\hat{\mathbf{y}}_{t+1,i}\|_2 \geq D\}} \right).$$
- 7: Meta-algorithm updates weight by $p_{t+1,i} \propto \exp(-\varepsilon_{t+1} \sum_{s=1}^t \langle \nabla g_s(\mathbf{y}_s), \mathbf{y}_{s,i} \rangle)$, $i \in [N]$.
- 8: Compute $\mathbf{y}_{t+1} = \sum_{i=1}^N p_{t+1,i} \mathbf{y}_{t+1,i}$.
- 9: Submit $\mathbf{x}_{t+1} = \Pi_{\mathcal{X}}[\mathbf{y}_{t+1}]$. \triangleright the only projection onto feasible domain \mathcal{X} per round
- 10: **end for**

For smooth and non-negative functions, the Sword++ algorithm (Zhao et al., 2024) achieves an $\mathcal{O}(\sqrt{(F_T^u + P_T)(1 + P_T)})$ small-loss dynamic regret, which requires one gradient and one function value per iteration.² However, notice that the surrogate loss $g_t(\cdot)$ in Eq. (6) is neither smooth nor non-negative, which hinders the application of our reduction to their method. Fortunately, owing to the benign property of $\|\nabla g_t(\mathbf{y}_t)\|_2 \leq \|\nabla f_t(\mathbf{x}_t)\|_2$ (see Theorem 1), we can still deploy the reduction via an improved analysis and obtain a projection-efficient algorithm with the same small-loss bound.

Theorem 3. *Set the step size pool as $\mathcal{H} = \{\eta_i = 2^{i-1} \sqrt{5D^2/(1 + 8LGDT)} \mid i \in [N]\}$ with $N = \lceil 2^{-1} \log_2((5D^2 + 2D^2T)(1 + 8LGDT)/(5D^2)) \rceil + 1$ and the learning rate of the meta-algorithm as $\varepsilon_t = \sqrt{(\ln N)/(1 + D^2 \sum_{s=1}^{t-1} \|\nabla g_s(\mathbf{y}_s)\|_2^2)}$. Under Assumptions 1, 2, and 3, our algorithm requires one projection onto \mathcal{X} per round and enjoys the following dynamic regret:*

$$\sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}_t) \leq \mathcal{O} \left(\sqrt{(F_T^u + P_T)(1 + P_T)} \right),$$

where $F_T^u = \sum_{t=1}^T f_t(\mathbf{u}_t)$ is the cumulative loss of the dynamic comparators.

3. Adaptive Regret Minimization

In this section, we present our efficient methods to minimize adaptive regret. First, we briefly review existing methods in Section 3.1, and then present our efficient methods to reduce the projection complexity of adaptive regret methods in Section 3.2.

2. The Sword++ algorithm is primarily designed to attain gradient-variation dynamic regret, incorporating advanced components like a correction term and optimism into its algorithmic design (Zhao et al., 2024). Nonetheless, it can be verified that when seeking a small-loss bound only, the algorithm's complexity can be reduced by omitting both the correction term and optimism.

3.1 A Brief Review of Adaptive Regret Minimization

Adaptive regret minimization ensures that the online learner is competitive with a fixed decision across every contiguous interval $I \subseteq [T]$. Typically, an online algorithm to achieve this consists of three components:

- (i) base-algorithm: an online algorithm attaining low (static) regret in a given interval;
- (ii) scheduling: a series of intervals that can cover the entire time horizon $[T]$, which might overlap. Each interval is associated with a base-learner whose goal is to minimize static regret over the duration of that interval (from its start to end);
- (iii) meta-algorithm: a combining algorithm that can track the best base-learner on the fly.

By dividing the entire algorithm into these three main components, it becomes more convenient to compare various algorithms and highlight the effectiveness of individual components.

For the worst-case bound, the best known result is the $\mathcal{O}(\sqrt{|I| \log T})$ adaptive regret achieved by the CBCE algorithm (Jun et al., 2017). We omit its details but mention that CBCE requires multiple gradients at each round. Wang et al. (2018) improve CBCE by employing the linearized loss for updating both meta-algorithm and base-algorithm. This revision allows it to require only one gradient per iteration while maintaining the same adaptive regret. Furthermore, this improved CBCE algorithm evaluates the function value only once per iteration. As a result, our reduction can be directly applied, yielding a projection-efficient variant with the same adaptive regret.

Now, we focus on the more challenging case of small-loss adaptive regret. The best known result is the $\mathcal{O}(\sqrt{F_I \log F_I \log F_T})$ adaptive regret for any interval $I = [r, s] \subseteq [T]$ obtained by the SACS algorithm (Zhang et al., 2019), where $F_I = \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=r}^s f_t(\mathbf{x})$ and $F_T = \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T f_t(\mathbf{x})$. However, SACS does not satisfy our reduction requirements. This is because it requires N gradient queries (specifically, $\nabla f_t(\mathbf{x}_{t,i})$ for $i \in [N]$) and $N + 1$ function evaluations (specifically, $f_t(\mathbf{x}_{t,i})$ for $i \in [N]$, and $f_t(\mathbf{x}_t)$) at round $t \in [T]$. Here, N denotes the number of active base-learners, and $\mathbf{x}_{t,i}$ denotes the local decision returned by the i -th base-learner. Therefore, we have to modify the algorithm to fit our purpose.

To this end, we need to review the construction of the SACS algorithm. First, SACS uses the scale-free online gradient descent (SOGD) (Orabona and Pál, 2018) as its base-algorithm, which ensures a small-loss regret in a given interval. Second, SACS employs AdaNormalHedge (Luo and Schapire, 2015) as the meta-algorithm, which supports the sleeping expert setup and also the small-loss regret. Finally, SACS introduces a novel scheduling strategy called the *problem-dependent geometric covering intervals*. This ensures that the number of maintained base-learners also depends on small-loss quantities. Owing to these designs, SACS can achieve a fully problem-dependent adaptive regret of order $\mathcal{O}(\sqrt{F_I \log F_I \log F_T})$, which scales according to the cumulative loss of comparators.

However, there are also some pitfalls. SACS also suffers from an $\mathcal{O}(\log T)$ projection complexity in the worst case due to its two-layer structure. Further, it can be observed that SACS only attains an $\mathcal{O}(\sqrt{|I| \log |I| \log T})$ bound in the worst case, which exhibits an $\mathcal{O}(\sqrt{\log |I|})$ gap compared with the best known result of $\mathcal{O}(\sqrt{|I| \log T})$ (Jun et al., 2017).

In the next subsection, we will present an efficient algorithm for small-loss adaptive regret minimization, which resolves the above two issues simultaneously.




Figure 1: Geometric Covering Intervals (Zhang et al., 2019, Figure 2).

3.2 Efficient Algorithms for Adaptive Regret

Since SACS involves multiple gradient and function queries in all its three components, we need to make modifications to achieve an algorithm that attains the same small-loss adaptive regret while demanding merely one gradient query and function evaluation per iteration. Once equipped with such an algorithm, we can deploy our reduction scheme to obtain an efficient method with 1 projection complexity.

The overall procedures of our proposed algorithm are summarized in Algorithm 2. Below, we will present the details of the three components. In particular, we will elucidate the scheduling design, i.e., the construction of covering intervals, which is paramount in achieving a fully problem-dependent bound (even improving upon the previously best known result of Zhang et al. (2019)) and reducing the number of gradient and function evaluations needed.

By the reduction mechanism, it is noticeable that we only need to consider the input online functions as surrogate loss $\{g_t\}_{t=1}^T$, where $g_t : \mathcal{Y} \mapsto \mathbb{R}$ is defined in Eq. (6).

Scheduling. Adaptive regret examines every contiguous interval $I \subseteq [T]$, which demands a rapid adaptation to potential environment changes. A natural way to construct the scheduling is to initiate a base-learner at each round and enable her to make predictions till the end of the whole time horizon (Hazan and Seshadhri, 2007). While this approach can effectively maintain many base-learners to handle the non-stationarity, it is computationally expensive due to maintaining $\mathcal{O}(t)$ base-learners at round t . For enhanced efficiency, alternative scheduling is proposed to set the length of each learner’s active time in a geometric manner (Hazan and Seshadhri, 2007; Daniely et al., 2015). To elucidate the design, we introduce two concepts: the *unit interval* and the *marker*. The unit intervals partition the time horizon $[T]$, of which the adaptive algorithm chooses geometric many to construct the active time as illustrated later. The markers denote the starting and ending time stamps of the unit intervals. Formally, the i -th unit interval is represented by $[s_i, s_{i+1} - 1]$ with the time stamps s_1, \dots, s_M determining the intervals (referred to as markers). Notice that, these unit intervals are disjoint and consecutive. The second and third rows of Figure 1 provide an illustrative example of unit intervals and markers. Based on the two concepts, we can then illustrate how to generate *geometric covering intervals*, also referred to as the scheduling above. The distinction between unit intervals and geometric intervals merits emphasis. The geometric covering intervals are the active time of each base-learner, and our adaptive algorithm manages the prediction period of the base-learners according to these intervals. The unit intervals partition the time horizon, and the covering intervals are generated based on these unit intervals by selecting

a geometric number of them. The adaptive algorithm initializes a base-learner once at the beginning time of each unit interval, say at time stamp s_m , and determines the active time based on the indexes of the markers,

$$[s_{i \cdot 2^k}, s_{(i+1) \cdot 2^k} - 1], \quad (9)$$

where $m = i \cdot 2^k, k \in \{0\} \cup \mathbb{N}$ and k is the largest number made the factorization. The setting indicates that the base-learner initialized at time s_m remains active across $2^k - 1$ unit intervals, leading to an exponentially number of unit intervals. This scheduling diversifies multiple base-learners to capture the non-stationarity across different time durations, and meanwhile ensures that at most $\mathcal{O}(\log T)$ base-learners are maintained per round.

Conventionally, previous studies set $s_t = t$ (Hazan and Seshadhri, 2007; Daniely et al., 2015), where the increasing of markers coincidences with one of the time stamps. This leads to the standard geometric covering intervals \mathcal{C} as shown in Figure 1, formally defined below,

$$\mathcal{C} = \bigcup_{k \in \mathbb{N} \cup \{0\}} \mathcal{C}_k, \text{ where } \mathcal{C}_k = \left\{ [i \cdot 2^k, (i+1) \cdot 2^k - 1] \mid i \text{ is odd} \right\} \text{ for } k \in \mathbb{N} \cup \{0\}, \quad (10)$$

which are derived from Eq. (9) by registering the marker at each round. However, this manner of initializing base-learners is *problem-independent*, which introduces a $\sqrt{\log T}$ factor to the adaptive regret bound of order $\mathcal{O}(\sqrt{|I| \log T})$, given that the meta-algorithm requires combining $\mathcal{O}(T)$ (sleeping-)experts in total. To address the issue, Zhang et al. (2019) propose the *problem-dependent* scheduling aiming to achieve a small-loss adaptive regret. In this paper, we further refine their design in order to reduce the gradient and function value queries at each round, which also helps improve their small-loss bound slightly.

Specifically, we aim to develop a fully data-dependent adaptive regret bound of order $\mathcal{O}(\sqrt{F_I \log F_I \log F_T})$ with 1 projection complexity. This bound replaces the $\mathcal{O}(\sqrt{\log T})$ factor by $\mathcal{O}(\sqrt{\log F_T \log F_I})$, which under specific beneficial conditions is considerably smaller. We postpone the discussion of the additional factor $\mathcal{O}(\sqrt{\log F_I})$ in Remark 1. To achieve this result, the SACS algorithm (Zhang et al., 2019) modifies the generation mechanism of the markers, which *registers a marker only when the cumulative loss exceeds a pre-defined fixed threshold*, instead of setting $s_t = t$. As a result, the number of active base-learners relates to the small-loss quantity, leading the overall algorithm to achieve a fully problem-dependent adaptive regret. The underlying rationale for this design stems from the potential inefficiency of the prior scheduling mechanism (which sets $s_t = t$ and initializes a base-learner at each round), since it is unnecessary to initialize a fresh new base-learner when the environment is relatively stable, or more precisely, when the cumulative loss is not large enough. Figure 2 provides an illustrative example of problem-dependent geometric covering intervals. Therefore, we are in the position to determine the threshold. SACS sets the threshold by monitoring the cumulative loss of the latest base-learner $f_t(\mathbf{x}_{t, i^\dagger})$ with i^\dagger being the latest base-learner's index, but this will introduce an additional function evaluation in addition to $f_t(\mathbf{x}_t)$ at each round. To avoid the limitation, we design the following two important improvements:

- we register markers and start a new base-learner according to the cumulative loss of *final decisions*, i.e., $\{f_t(\mathbf{x}_t)\}_{t=1}^T$, bypassing the additional function evaluation;
- we introduce a sequence of *time-varying thresholds* with a careful design, instead of using a fixed threshold over the time horizon.




Figure 2: Problem-dependent Geometric Covering Intervals (Zhang et al., 2019, Figure 4).

This configuration of covering intervals realizes the condition of one function evaluation per round. Additionally, the new design of the thresholds mechanism is crucial to ensure that the small-loss bound can simultaneously recover the best known worst-case guarantee, which cannot be achieved by prior best small-loss adaptive regret bound (Zhang et al., 2019). More discussions are presented in Remark 1.

We are now ready to introduce our efficient algorithm for optimizing the small-loss adaptive regret. Let C_1, C_2, C_3, \dots denote the sequence of thresholds, determined by a threshold generating function $\mathcal{G}(\cdot) : \mathbb{N} \mapsto \mathbb{R}_+$, which we will specify later. Our problem-dependent geometric covering intervals are defined as follows. We initialize the setting by $s_1 = 1$. We set s_2 as the round when the cumulative loss of the overall algorithm (namely, $\sum_{s=1}^t f_s(\mathbf{x}_s)$) exceeds the threshold C_1 and then initialize a new instance of SOGD starting at this round. The process is repeated until the end of the online learning process. We thus generate a sequence of markers $\{s_1, s_2, \dots\}$. See the condition in Line 7, registration of markers in Line 9, and the overall updates in Lines 7 – 11 of Algorithm 2. Those markers specify the starting time (and the ending time) of base-learners and further we can construct the problem-dependent covering intervals as

$$\tilde{\mathcal{C}} = \bigcup_{k \in \mathbb{N} \cup \{0\}} \tilde{\mathcal{C}}_k, \text{ where } \tilde{\mathcal{C}}_k = \left\{ [s_{i \cdot 2^k}, s_{(i+1) \cdot 2^k} - 1] \mid i \text{ is odd} \right\} \text{ for all } k \in \mathbb{N} \cup \{0\}. \quad (11)$$

This construction matches exactly the method of selecting geometric many of the unit intervals as described in Eq. (9). It is worth noting that the problem-dependent covering intervals, unlike the standard ones which set $s_t = t$, are constructed using markers that cannot be specified with exact time stamps in advance. Instead, they are determined according to the learner's performance on the fly. However, this does not hinder the practice of our algorithm, as it activates or deactivates base-learners based on the marker indices it maintains.

Base-algorithm. We employ SOGD as the base-algorithm, running with a *linearized loss* $\langle \nabla g_t(\mathbf{y}_t), \mathbf{y} \rangle$ over the surrogate domain \mathcal{Y} . Denote by A_t the set of active base-learners' indices, then the base-learner \mathcal{B}_i updates by

$$\mathbf{y}_{t+1,i} = \Pi_{\mathcal{Y}}[\mathbf{y}_{t,i} - \eta_{t,i} \nabla g_t(\mathbf{y}_t)],$$

with $\eta_{t,i} = D / \sqrt{(\delta + \sum_{s=\tau_i}^t \|\nabla g_s(\mathbf{y}_s)\|_2^2)}$, where τ_i denotes the starting time of the base-learner $i \in A_t$. The projection onto \mathcal{Y} can be easily calculated by a simple rescaling if needed. Notably, owing to the convexity of the surrogate loss g_t , we can use the *same* gradient $\nabla g_t(\mathbf{y}_t)$ for all the base-learners at each round, ensuring one gradient query of $\nabla f_t(\mathbf{x}_t)$ at each round.

Algorithm 2 Efficient Algorithm for Problem-dependent Adaptive Regret

Input: threshold generating function $\mathcal{G}(\cdot) : \mathbb{N} \mapsto \mathbb{R}_+$.

- 1: Initialize total intervals $m = 1$, marker $s_1 = 1$, threshold $C_1 = \mathcal{G}(1)$; let \mathbf{x}_1 be any point in \mathcal{X} ; let A_t denote the set of indexes for the active base-learners at time t .
- 2: **for** $t = 1$ **to** T **do**
- 3: Receive the gradient information $\nabla f_t(\mathbf{x}_t)$.
- 4: Construct the surrogate loss $g_t : \mathcal{Y} \mapsto \mathbb{R}$ according to Eq. (6).
- 5: Compute the (sub-)gradient $\nabla g_t(\mathbf{y}_t)$ according to Lemma 1.
- 6: Compute $L_t = L_{t-1} + f_t(\mathbf{x}_t)$.
- % constructing Problem-dependent Geometric Covers(PGC)
- 7: **if** $L_t > C_m$ **then**
- 8: Set $L_t = 0$, remove base-learners \mathcal{B}_k whose deactivating time stamp is before the registration of $(m+1)$ -th marker.
- 9: Set $m \leftarrow m+1$, register marker $s_m \leftarrow t$, update threshold $C_m = \mathcal{G}(m)$.
- 10: Initialize a new base-learner whose active span is $[s_{n \cdot 2^k}, s_{(n+1) \cdot 2^k} - 1]$ where $m = n \cdot 2^k, k \in \{0\} \cup \mathbb{N}$ and k is the largest number made the factorization.
- 11: Set $\gamma_m = \ln(1 + 2m)$, $w_{t,m} = 1$, $\eta_{t,m} = \min\{1/2, \sqrt{\gamma_m}\}$ for the meta-algorithm.
- 12: **end if**
- 13: Send $\nabla g_t(\mathbf{y}_t)$ to all base-learners and obtain local predictions $\mathbf{y}_{t+1,i}$ for $i \in A_t$.
- 14: Meta-algorithm updates weight $\mathbf{p}_{t+1} \in \Delta_{|A_{t+1}|}$ by Eq. (12), Eq. (13), and Eq. (14)
- 15: Compute $\mathbf{y}_{t+1} = \sum_{i \in A_{t+1}} p_{t+1,i} \mathbf{y}_{t+1,i}$.
- 16: Submit $\mathbf{x}_{t+1} = \Pi_{\mathcal{X}}[\mathbf{y}_{t+1}]$. \triangleright the only projection onto feasible domain \mathcal{X} per round
- 17: **end for**

Meta-algorithm. SACS uses AdaNormalHedge (Luo and Schapire, 2015) as its meta-algorithm, but this choice is not suitable for us. To ensure only one projection per iteration, we cannot rely on multiple function values $\{g_t(\mathbf{y}_{t,i})\}_{i=1}^{|A_t|}$ for the meta-algorithm to evaluate the loss. Instead, we can only use the *linearized* loss value, namely, $\{\langle \nabla g_t(\mathbf{y}_t), \mathbf{y}_{t,i} \rangle\}_{i=1}^{|A_t|}$ in the meta update. The small-loss bound in SACS crucially depends on access to the original function values, which are unavailable here. Technically, with only linearized losses, it is difficult to establish a *squared* gradient-norm bound and then convert it to the small-loss bound due to the *first-order* regret bound of AdaNormalHedge. This key observation motivates us to instead use the Adapt-ML-Prod algorithm (Gaillard et al., 2014) as our meta-algorithm. The main advantage is that it enjoys a *second-order regret* while also supporting the sleeping expert setup. Adapt-ML-Prod maintains multiple learning rates $\boldsymbol{\eta}_{t+1} \in [0, 1]^{|A_{t+1}|}$ and an intermediate weight vector $\mathbf{w}_{t+1} \in \mathbb{R}^{|A_{t+1}|}$, , updated as follows: for any active base-learner $i \in A_{t+1}$,

$$\eta_{t+1,i} = \min \left\{ \frac{1}{2}, \sqrt{\frac{\gamma_i}{1 + \sum_{k=s_i}^t (\hat{\ell}_k - \ell_{k,i})^2}} \right\}, \quad w_{t+1,i} = \left(w_{t,i} (1 + \eta_{t,i} (\hat{\ell}_t - \ell_{t,i})) \right)^{\frac{\eta_{t+1,i}}{\eta_{t,i}}}, \quad (12)$$

where $\gamma_i = \ln(1 + 2i)$ is a certain scaling factor and the feedback loss is constructed in the following way, for $i \in A_t$, set

$$\hat{\ell}_t = \langle \nabla g_t(\mathbf{y}_t), \mathbf{y}_t \rangle / (2GD), \text{ and } \ell_{t,i} = \langle \nabla g_t(\mathbf{y}_t), \mathbf{y}_{t,i} \rangle / (2GD). \quad (13)$$

The final weight vector $\mathbf{p}_{t+1} \in \Delta_{|A_{t+1}|}$ is obtained by

$$p_{t+1,i} = \frac{w_{t+1,i} \cdot \eta_{t+1,i}}{\sum_{j \in A_{t+1}} w_{t+1,j} \cdot \eta_{t+1,j}}. \quad (14)$$

Notably, the meta update only uses one gradient at round t , namely, $\nabla g_t(\mathbf{y}_t)$.

Finally, we compute $\mathbf{y}_{t+1} = \sum_{i \in A_{t+1}} p_{t+1,i} \mathbf{y}_{t+1,i}$ as the overall prediction in the surrogate domain \mathcal{Y} and calculate $\mathbf{x}_{t+1} = \Pi_{\mathcal{X}}[\mathbf{y}_{t+1}]$ to ensure the feasibility. This is the only projection onto \mathcal{X} at each round. Algorithm 2 summarizes the main procedures of our efficient method for small-loss adaptive regret. Albeit with a similar two-layer structure as SACS, our algorithm exhibits salient differences in meta-algorithm and covering intervals. As a benefit, we can successfully deploy our reduction mechanism and make the overall algorithm project onto the feasible domain \mathcal{X} once per round, see Line 16. Our method retains the same small-loss adaptive regret as (Zhang et al., 2019).

Theorem 4. *Under Assumptions 1–3, setting the threshold generating function $\mathcal{G} : \mathbb{N} \mapsto \mathbb{R}$,*

$$\mathcal{G}(m) = (54GD + 168D^2L) \ln(1 + 2m) + 168D^2L\mu_T^2 + 18GD\mu_T + 6D\sqrt{\delta} + 672D^2L, \quad (15)$$

where $\mu_T = \ln(1 + (1 + \ln(1 + T))/(2e))$ and thus $\mathcal{G}(m) = \mathcal{O}(\log m)$, Algorithm 2 requires only one projection onto the feasible domain \mathcal{X} per round and enjoys the following small-loss adaptive regret for any interval $I = [r, s] \subseteq [T]$:

$$\sum_{t=r}^s f_t(\mathbf{x}_t) - \sum_{t=r}^s f_t(\mathbf{u}) \leq \mathcal{O}\left(\min\left\{\sqrt{F_I \log F_I \log F_T}, \sqrt{|I| \log T}\right\}\right), \quad (16)$$

where $F_I = \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=r}^s f_t(\mathbf{x})$ and $F_T = \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T f_t(\mathbf{x})$.

Remark 1. Note that the $\mathcal{O}(\sqrt{F_I \log F_I \log F_T})$ small-loss bound of Zhang et al. (2019) becomes $\mathcal{O}(\sqrt{|I| \log |I| \log T})$ in the worst case, looser than the $\mathcal{O}(\sqrt{|I| \log T})$ bound (Jun et al., 2017) by a factor of $\sqrt{\log |I|}$. We show that this limitation can be actually avoided by the new design of the threshold mechanism and a refined analysis, both of which are crucial for obtaining the additional $\mathcal{O}(\sqrt{|I| \log T})$ worst-case regret guarantee. Indeed, our result in (16) can *strictly* match the best known problem-independent result in the worst case.

4. Interval Dynamic Regret Minimization

As mentioned in Section 1.1, adaptive regret and dynamic regret are not directly comparable under the general OCO settings. Therefore, researchers consider optimizing them *simultaneously* using a more stringent measure called *interval dynamic regret*, which competes with any changing comparator sequence over *any* interval (Zhang et al., 2020; Cutkosky, 2020).

4.1 A Brief Review of Interval Dynamic Regret Minimization

Zhang et al. (2020) propose the AOA algorithm to optimize the interval dynamic regret and obtain a worst-case bound of $\sum_{t=r}^s f_t(\mathbf{x}_t) - \sum_{t=r}^s f_t(\mathbf{u}_t) \leq \mathcal{O}(\sqrt{|I|(\log T + P_I)})$ for any interval $I = [r, s] \subseteq [T]$, where $P_I = \sum_{t=r+1}^s \|\mathbf{u}_t - \mathbf{u}_{t-1}\|_2$ denotes the path length. The overall structure of AOA is similar to the adaptive algorithm framework described in

Section 3.2. The major difference is that AOA uses Ader (Zhang et al., 2018a), a two-layer structure algorithm, as the base algorithm to minimize dynamic regret over a given interval. As a result, the projection complexity of their method is $\mathcal{O}(\log^2 T)$, since each base-learner requires $\mathcal{O}(\log T)$ projections, and AOA employs $\mathcal{O}(\log T)$ such base-learners.

Cutkosky (2020) achieves an $\tilde{\mathcal{O}}(\sqrt{|I|(1 + P_I)})$ interval dynamic regret through a fundamentally different framework based on parameter-free online learning, where $\tilde{\mathcal{O}}(\cdot)$ hides the polynomial dependence on $\log T$ factors. In fact, Cutkosky (2020) derives a *gradient-norm bound* (scaling with $\sum_{t \in I} \|\nabla f_t(\mathbf{x}_t)\|^2$) for interval dynamic regret, which can be easily converted into a small-loss bound (scaling with $\sum_{t \in I} f_t(\mathbf{u}_t)$) when the online functions are non-negative and smooth. However, we note that, even in this case, their small-loss bound is not favorable enough, as the regret guarantee is not *fully* problem-dependent — it still exhibits polynomial dependence on $\log T$ terms.

4.2 Efficient Algorithms for Interval Dynamic Regret

In this part, we present efficient algorithms for optimizing interval dynamic regret. We begin by discussing the worst-case interval dynamic regret, which is relatively straightforward to achieve. Next, we focus on the more challenging small-loss regret bound.

Worst-case interval dynamic regret. Although AOA itself does not satisfy the condition of one gradient query and one function evaluation per round, it can be verified that simply using a linearized loss for AOA achieves the same bound while meeting the reduction requirement. Consequently, applying our reduction mechanism yields an efficient algorithm for optimizing the interval dynamic regret. This algorithm retains the same worst-case bound as AOA but significantly reduces the projection complexity from $\mathcal{O}(\log^2 T)$ to 1.

Small-loss interval dynamic regret. Regarding the small-loss type bound, the only relevant result in the literature is the gradient-norm bound presented by Cutkosky (2020), which can be further converted into a small-loss bound of order $\tilde{\mathcal{O}}(\sqrt{(F_I^{\mathbf{u}} + P_I)(1 + P_I)})$ when the online functions are non-negative and smooth, where $F_I^{\mathbf{u}} = \sum_{t \in I} f_t(\mathbf{u}_t)$ and $P_I = \sum_{t \in I} \|\mathbf{u}_t - \mathbf{u}_{t-1}\|_2$. However, it is only a *partially* problem-dependent bound, as the $\tilde{\mathcal{O}}(\cdot)$ notation still hides polynomial dependence on $\log T$ factors.

In the following, we present the first *fully* problem-dependent small-loss bound for optimizing the interval dynamic regret, building on the AOA algorithm (Zhang et al., 2020). Importantly, our algorithm requires only a single projection onto the feasible domain \mathcal{X} per round.³ To achieve this, we employ the same structure as AOA, but with several modifications. Specifically, we use a linearized loss for the base-learners and Adapt-ML-Prod as the meta-algorithm. Additional effort is required to construct the problem-dependent geometric covers. Algorithm 3 summarizes the main procedures. The structure of this algorithm is similar to our algorithm for small-loss adaptive regret (see Algorithm 2), with the key difference that we replace the base-algorithm from SOGD with our efficient method for dynamic regret (i.e., Algorithm 1). This change is necessary to manage the *dual uncertainties* inherent in interval dynamic regret minimization, arising from both the unknown interval and

3. Our reduction mechanism can also be applied to reduce the projection complexity of the method proposed by Cutkosky (2020), but the variant still exhibits a less favorable regret and higher computational complexity compared to the one presented here. Further details can be found in Remark 3 and Remark 4.

Algorithm 3 Efficient Algorithm for Problem-dependent Interval Dynamic Regret

Input: step size pool $\mathcal{H} = \{\eta_1, \dots, \eta_N\}$, learning rate setting $\varepsilon_{i,t}$ for the i -th base-learner, and the threshold generating function $\mathcal{G}(\cdot) : \mathbb{N} \mapsto \mathbb{R}_+$.

- 1: Set total intervals $m = 1$, marker $s_1 = 1$, threshold $C_1 = \mathcal{G}(1)$; let $\mathbf{x}_1 \in \mathcal{X}$ be any point; let A_t denote the set of indexes for the active base-learners at time t .
- 2: **for** $t = 1$ **to** T **do**
- 3: Receive the gradient information $\nabla f_t(\mathbf{x}_t)$.
- 4: Construct the surrogate loss $g_t : \mathcal{Y} \mapsto \mathbb{R}$ according to Eq. (6).
- 5: Compute the (sub-)gradient $\nabla g_t(\mathbf{y}_t)$ according to Lemma 1.
- 6: Compute $L_t = L_{t-1} + f_t(\mathbf{x}_t)$.
- % constructing Problem-dependent Geometric Covers(PGC)**
- 7: **if** $L_t > C_m$ **then**
- 8: Set $L_t = 0$, remove base-learners \mathcal{B}_k whose deactivating time stamp is before the registration of $(m+1)$ -th marker.
- 9: Set $m \leftarrow m+1$, register marker $s_m \leftarrow t$, update threshold $C_m = \mathcal{G}(m)$.
- 10: Initialize a new base-learner according to Eq. (17) and Eq. (18), whose active span is $[s_{n \cdot 2^k}, s_{(n+1) \cdot 2^k} - 1]$ where $m = n \cdot 2^k, k \in \{0\} \cup \mathbb{N}$ and k is the largest number made the factorization.
- 11: Set the required inputs with step size pool \mathcal{H} , learning rate for Hedge used in base-algorithm as $\varepsilon_{i,t}$ with $i = m$.
- 12: Set $\gamma_m = \ln(1 + 2m)$, $w_{t,m} = 1$, $\eta_{t,m} = \min\{1/2, \sqrt{\gamma_m}\}$ for the meta-algorithm.
- 13: **end if**
- 14: Send $\nabla g_t(\mathbf{y}_t)$ to all base-learners and obtain local predictions $\mathbf{y}_{t+1,i}$ for $i \in A_t$.
- 15: Meta-algorithm updates weight $\mathbf{p}_{t+1} \in \Delta_{|A_{t+1}|}$ by Eq. (12), Eq. (13), and Eq. (14).
- 16: Compute $\mathbf{y}_{t+1} = \sum_{i \in A_{t+1}} p_{t+1,i} \mathbf{y}_{t+1,i}$.
- 17: Submit $\mathbf{x}_{t+1} = \Pi_{\mathcal{X}}[\mathbf{y}_{t+1}]$. \triangleright the only projection onto feasible domain \mathcal{X} per round
- 18: **end for**

the unknown comparator sequence. Intuitively, the design of this interval dynamic algorithm can be thought of as a *three-layer structure*, because, in addition to using Adapt-ML-Prod to combine the base-learners' decisions, each base-learner also generates her own decision by combining several maintained subroutines (OGD) with a meta-algorithm (Hedge with time-varying learning rates). The details are as follows.

Scheduling and Meta-algorithm. To attain the small-loss interval dynamic regret, we employ the same scheduling mechanism and meta-algorithm as Algorithm 2 (which is designed for small-loss adaptive regret), namely, we use the problem-dependent geometric covering intervals determined by a sequence of time-varying thresholds and use the Adapt-ML-Prod as the meta-algorithm. One can refer to Section 3.2 for detailed elaborations.

Base-algorithm. As mentioned, we now employ our efficient method for dynamic regret minimization (i.e., Algorithm 1) as the base-algorithm, which consists two layers. We specify the update procedure of the i -th base-learner. At round $t+1$, the base-learner will submit

the decision as $\mathbf{y}_{t+1,i} = \sum_{j=1}^N p'_{t+1,i,j} \mathbf{y}_{t+1,i,j}$, where

$$\text{base:base-level: } \mathbf{y}_{t+1,i,j} = \Pi_{\mathcal{Y}} [\mathbf{y}_{t+1,i,j} - \eta_j \nabla g_t(\mathbf{y}_t)], j \in [N], \quad (17)$$

$$\text{base:meta-level: } p'_{t+1,i,j} \propto \exp \left(-\varepsilon_{i,t+1} \sum_{s=i_{\text{start}}}^t \langle \nabla g_s(\mathbf{y}_s), \mathbf{y}_{s,i,j} \rangle \right), j \in [N], \quad (18)$$

We use the prefix “base:” to indicate updates performed within the base-algorithm utilized by Algorithm 3. The local decision $\mathbf{y}_{t+1,i,j}$ is returned from the j -th base:base-learner in the base-algorithm, N denotes the number of base:base-learners, i_{start} denotes the starting time stamp for the i -th base-learner and $p_{t+1,i,j}$ denotes its corresponding combination weight. Note that these base-level updates are conducted within the domain \mathcal{Y} , where the algorithm can benefit from a rapid projection. Moreover, it is also worth mentioning the choice of step size η_j . Ideally, it should be scaled with the length of the active interval $|I|$ of the associated base-learner, nevertheless, we actually set it based on T in Theorem 5 due to the problem-dependent covering intervals. Technically, since the active time of each base-learner is determined on the fly, we can only set potentially over-estimated learning rates to ensure that the dynamic regret is guaranteed whenever the base-learner is deactivated. As shown in Lemma 8, this setup ensures that each base-learner enjoys an anytime dynamic regret bound.

With these components in place, the following theorem establishes the small-loss interval dynamic regret of Algorithm 3.

Theorem 5. *Under Assumptions 1–3, employing the step size pool for base:base-algorithm (defined at Eq. (17)) as $\mathcal{H} = \{\eta_j = 2^{j-1} \sqrt{5D^2/(1+8LGD T)} \mid j \in [N]\}$ with $N = \lceil 2^{-1} \log_2((5D^2 + 2D^2 T)(1+8LGD T)/(5D^2)) \rceil + 1$, setting the meta-algorithm’s learning rate of the i -th base-learner (defined at Eq. (18)) as $\varepsilon_{i,t} = \sqrt{(\ln N)/(1 + D^2 \sum_{s=i_{\text{start}}}^{t-1} \|\nabla g_s(\mathbf{y}_s)\|_2^2)}$ with i_{start} denoting its starting time stamp, setting the threshold generating function*

$$\begin{aligned} \mathcal{G}(m) = & 7L \left(12D\sqrt{\ln(1+2m)} + 4D\mu_T + 6D\sqrt{\ln N} \right)^2 \\ & + 54GD \ln(1+2m) + 18GD\mu_T + \frac{3(6+G^2D^2)\sqrt{\ln N}}{2} + (630L+23)D^2 + 9, \end{aligned} \quad (19)$$

where $\mu_T = \ln(1+(1+\ln(1+T))/(2e))$ and thus $\mathcal{G}(m) = \mathcal{O}(\log m)$, Algorithm 3 requires only 1 projection onto \mathcal{X} per round and ensures that the interval dynamic regret $\sum_{t=r}^s f_t(\mathbf{x}_t) - \sum_{t=r}^s f_t(\mathbf{u}_t)$ is upper bounded by

$$\mathcal{O} \left(\sqrt{(F_I^{\mathbf{u}} + P_I) \cdot (P_I + \log(\min_{\mathbf{u}_{1:T}} \{F_T^{\mathbf{u}} + P_T\})) \cdot \log(\min_{\mathbf{u}_{r:s}} \{F_I^{\mathbf{u}} + P_I\})} \right), \quad (20)$$

which holds for any interval $I = [r, s] \subseteq [T]$ and any comparator $\mathbf{u}_1, \dots, \mathbf{u}_T \in \mathcal{X}$. Here, $P_I = \sum_{t=r+1}^s \|\mathbf{u}_t - \mathbf{u}_{t-1}\|_2$ denotes the path length, $F_I^{\mathbf{u}} = \sum_{t=r}^s f_t(\mathbf{u}_t)$ denotes the comparators’ cumulative loss, $\min_{\mathbf{u}_{r:s}} \{F_I^{\mathbf{u}} + P_I\} = \min_{\mathbf{u}_r, \dots, \mathbf{u}_s} \{\sum_{t=r}^s f_t(\mathbf{u}_t) + \sum_{t=r+1}^s \|\mathbf{u}_t - \mathbf{u}_{t-1}\|_2\}$, and we abbreviate $P_{[1:T]}$ and $F_{[1:T]}$ by P_T and F_T respectively.

The proof can be found in Appendix D. Given Theorem 3 and Theorem 4, a natural conjecture of interval dynamic regret bound might be $\mathcal{O}(\sqrt{(F_I^{\mathbf{u}} + P_I) \cdot (P_I + \log F_T^{\mathbf{u}}) \cdot \log F_I^{\mathbf{u}}})$.

Nonetheless, our analysis reveals that in the context of interval dynamic regret, $F_I^{\mathbf{u}} + P_I$ should actually be treated as a combined quantity and estimated together when designing the threshold mechanism. A precise analysis can be found in Lemma 10 and Lemma 11.

The obtained result is versatile. When competing with a fixed comparator, the path length $P_I = 0$ and interval dynamic regret (20) becomes $\mathcal{O}(\sqrt{F_I \log F_I \log F_T})$, recovering the small-loss adaptive regret exhibits in Theorem 4. Moreover, when considering the entire time horizon with $I = [1, T]$, the regret bound becomes $\tilde{\mathcal{O}}(\sqrt{(F_T^{\mathbf{u}} + P_T)(1 + P_T)})$, which nearly matches the $\mathcal{O}(\sqrt{(F_T^{\mathbf{u}} + P_T)(1 + P_T)})$ dynamic regret in Theorem 3 up to logarithmic factors in $\min_{\mathbf{u}_{1:T}} \{F_T^{\mathbf{u}} + P_T\}$. We note that a similar logarithmic gap of $\mathcal{O}(\log T)$ also exists in the study of the problem-independent interval dynamic regret (Zhang et al., 2020, Theorem 5).

Algorithm 3 requires only one gradient query, one function value, and one projection operation, improving the oracle query complexity over previous method (Zhang et al., 2020) for interval dynamic regret minimization. Moreover, Algorithm 3 needs to maintain $\mathcal{O}(\log T \cdot \log(\min_{\mathbf{u}_{1:T}} \{F_T^{\mathbf{u}} + P_T\}))$ instances of online gradient descent, where $\log(\min_{\mathbf{u}_{1:T}} \{F_T^{\mathbf{u}} + P_T\})$ can be a constant if the environments are favorable. This approach enhances the efficiency over that of Zhang et al. (2020), which requires $\mathcal{O}(\log^2 T)$ instances. Compared to Algorithm 2, which is designed for minimizing adaptive regret and requires $\mathcal{O}(\log F_T)$ instances, we demonstrate that $\min_{\mathbf{u}_{1:T}} \{F_T^{\mathbf{u}} + P_T\} \leq F_T$, since the choices of $\mathbf{u}_1 = \dots = \mathbf{u}_T = \mathbf{u}^* \in \arg \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T f_t(\mathbf{x})$ will lead to $F_T^{\mathbf{u}} + P_T = F_T$, which may not be the optimal solution to $\min_{\mathbf{u}_{1:T}} \{F_T^{\mathbf{u}} + P_T\}$. Therefore, Algorithm 3 can even have the computational advantage over the two-layer structure algorithm if $F_T \gg \min_{\mathbf{u}_{1:T}} \{F_T^{\mathbf{u}} + P_T\}$, which may occur in non-stationary environments, e.g., where the best decision in each round \mathbf{x}_t^* is drifting slowly.

Cutkosky (2020) achieves an interval dynamic regret bound with techniques from parameter-free online learning. It is worthwhile to examine the subtle differences between our findings and those of Cutkosky (2020); the following remarks provide this comparison.

Remark 2. Algorithm 3 actually enjoys the same gradient-norm bound as Cutkosky (2020, Theorem 7) (in fact even stronger by logarithmic factors in T) under Assumptions 1 and 2 (without requiring the smoothness assumption). More specifically, our analysis is conducted in terms of $\|\nabla g_t(\mathbf{y}_t)\|_2$ (see Eq. (65) and Eq. (66)), which can then be related to $\|\nabla f_t(\mathbf{x}_t)\|_2$ by Lemma 1. Following the same convention of Cutkosky (2020), who uses the $\tilde{\mathcal{O}}(\cdot)$ -notation to hide the logarithmic dependence in T , Algorithm 3 can obtain an $\tilde{\mathcal{O}}(\sqrt{(G_I + P_I)(1 + P_I)})$ interval dynamic regret for any interval $I = [r, s] \subseteq [T]$ without any modification on the learning rates, where $P_I = \sum_{t=r+1}^s \|\mathbf{u}_t - \mathbf{u}_{t-1}\|_2$ and $G_I = \sum_{t=r}^s \|\nabla f_t(\mathbf{x}_t)\|_2^2$.

Remark 3. The interval dynamic regret bound obtained by Cutkosky (2020) essentially implies a small-loss bound $\tilde{\mathcal{O}}(\sqrt{(F_I^{\mathbf{u}} + P_I)(1 + P_I)})$ when assuming the smoothness of loss functions. However, the $\tilde{\mathcal{O}}(\cdot)$ notation hides problem-independent factors such as $\log T$. In contrast, our result in Theorem 5 is fully problem-dependent, replacing the quantity such as $\log T$ by $\log(\min_{\mathbf{u}_{1:T}} \{F_T^{\mathbf{u}} + P_T\})$, through our techniques of constructing problem-dependent covering intervals and corresponding threshold designs. These problem-dependent quantities can be much smaller under nice environments, and it remains unclear how to adapt our techniques to the method of Cutkosky (2020) due to a different framework.

Remark 4. The proposed Algorithm 3 requires only one projection onto the original domain, whereas the method by Cutkosky (2020) requires multiple projections onto a more complex,

lifted domain per round, raising concerns about computational efficiency. In Appendix D.1, we present how to apply our reduction mechanism to the method of Cutkosky (2020), replacing most of the potentially costly projections onto the lifted domains by projections onto cylinders. However, one projection onto the complex domain remains necessary.

5. Applications

In this section, we provide two applications of our proposed reduction mechanism for non-stationary online learning, including minimizing the *dynamic regret of online non-stochastic control* and minimizing the *adaptive regret of online principal component analysis*.

It is important to note that both problems are adaptations of the standard online learning settings with essential modifications. Specifically, there are three key points worth mentioning: (i) both problems operate over the matrix space instead of the vector space studied in previous sections; (ii) for online non-stochastic control, our reduction mechanism is further enhanced to account for the *switching cost* of algorithmic decisions (Anava et al., 2015; Agarwal et al., 2019a), a crucial characteristic of this decision-theoretic problem; (iii) for online principal component analysis, we actually present the first *strongly* adaptive regret result (still with one projection per round), improving the regret bound upon the previously best known result (Yuan and Lamperski, 2019) that only achieves weakly adaptive regret.

5.1 Online Non-Stochastic Control

In this part, we apply our reduction mechanism to an important online decision-making problem, online non-stochastic control, which attracts much attention these years in online learning and control theory community (Agarwal et al., 2019b; Foster and Simchowitz, 2020; Hazan et al., 2020; Gradu et al., 2020; Zhao et al., 2023; Hazan and Singh, 2022).

5.1.1 PROBLEM FORMULATION

We focus on the online control of linear dynamical system (LDS) defined as $x_{t+1} = Ax_t + Bu_t + w_t$, where x_t is the state, u_t is the control, w_t is a disturbance to the system. The controller suffers cost $c_t(x_t, u_t)$ with convex function $c_t : \mathbb{R}^{d_x} \times \mathbb{R}^{d_u} \mapsto \mathbb{R}$. Throughout this subsection, we follow the convention of notations in the non-stochastic control community to use unbold fonts to represent vectors and matrices. In online non-stochastic control, since there are no statistical assumptions imposed on system disturbance w_t and additionally the cost function can be chosen adversarially. The adversarial nature of the control setting hinders us from precomputing the optimal policy, as is possible in classical control theory (Kalman, 1960), and therefore requires modern online learning techniques to tackle adversarial environments.

We adopt *dynamic policy regret* (Zhao et al., 2023) to benchmark the performance of the designed controller with a sequence of arbitrary *time-varying* controllers $\pi_1, \dots, \pi_T \in \Pi$,

$$\text{D-REG}(\pi_1, \dots, \pi_T) = \sum_{t=1}^T c_t(x_t, u_t) - \sum_{t=1}^T c_t(x_t^{\pi_t}, u_t^{\pi_t}). \quad (21)$$

For this problem, the pioneering work (Agarwal et al., 2019a) investigates the static regret as (21) with $\pi_1, \dots, \pi_T \in \arg \min_{\pi \in \Pi} \sum_{t=1}^T c_t(x_t^\pi, u_t^\pi)$. The authors propose a gradient-based

controller with $\tilde{\mathcal{O}}(\sqrt{T})$ static regret. Specifically, they propose to employ the Disturbance-Action Controller (DAC) policy class $\pi(K, M)$, which is parametrized by a fixed matrix $K \in \mathbb{R}^{d_u \times d_x}$ and parameters matrix tuple $M = (M^{[1]}, \dots, M^{[H]}) \in (\mathbb{R}^{d_u \times d_x})^H$ with a memory length H . At each round, DAC makes the decision as a linear map of the past disturbances with an offset linear controller $u_t = -Kx_t + \sum_{i=1}^H M^{[i]}w_{t-i}$. This parametrization makes the action as a linear function of the past disturbances and further can reduce the online non-stochastic control to online convex optimization with memory (OCO with Memory) (Anava et al., 2015) with truncated loss, and thus one can further apply techniques developed for OCO with memory to handle the non-stochastic control.

For online non-stochastic control, Zhao et al. (2023) propose an online control approach with an $\tilde{\mathcal{O}}(\sqrt{T(1 + P_T)})$ dynamic policy regret, where $P_T = \sum_{t=2}^T \|M_{t-1}^* - M_t^*\|_F$ denotes the cumulative variation of comparators. The algorithm leverage an online ensemble structure equipped with $\mathcal{O}(\log T)$ base-learners, which leads to an $\mathcal{O}(\log T)$ projection complexity. In the sequel, we will investigate the computational complexity of the projection operation. Our aim is to refine the method through our efficient reduction mechanism, obtaining an algorithm that retains the *same* regret guarantee while requiring *one* projection per round.

5.1.2 PROJECTION COMPUTATIONAL COMPLEXITY

Previous studies project the parameters matrix tuple M onto the following domain,

$$\mathcal{M} = \left\{ M \triangleq (M^{[1]}, \dots, M^{[H]}) \in \left(\mathbb{R}^{d_u \times d_x} \right)^H \mid \|M^{[i]}\|_{\text{op}} \leq c_i \right\}, \quad (22)$$

where $\|\cdot\|_{\text{op}}$ denotes the operator norm and c_i is some fixed constant. The projection can be done by projecting each matrix $M^{[i]}, i \in [H]$, onto spectral norm ball with radius c_i sequentially. For each $M^{[i]}$, the projection involves the dominant process of diagonalization and project the singular values onto ℓ_∞ ball. Thus, the projection computational complexity onto \mathcal{M} is in order of $\mathcal{O}(H \min\{d_u^2 d_x, d_u d_x^2\})$, which dominates the computational expense in the base-learners, given that the gradient descent step requires $\mathcal{O}(H \cdot d_u d_x)$ time only.

5.1.3 EFFICIENT REDUCTION

In this part, we apply our efficient reduction mechanism to Scream.Control (Zhao et al., 2023), which can improve the projection complexity and maintain the theoretical guarantee. The algorithm is summarized in Algorithm 4 and we introduce the main ingredients below.

To facilitate the efficient reduction, we design the following surrogate domain,

$$\mathcal{M}' = \left\{ M = (M^{[1]}, \dots, M^{[H]}) \in \left(\mathbb{R}^{d_u \times d_x} \right)^H \mid \|M^{[i]}\|_F \leq c_i \sqrt{d} \right\}, \quad (23)$$

as the replacement of the original one defined at Eq. (22), where we denote by $d = \min\{d_x, d_u\}$. Notice that Scream.Control already satisfies the reduction requirements of querying only one function value and one gradient value per round (Zhao et al., 2023). Also, this algorithm utilizes linearized cost function to perform update for meta-learner and base-learners, enabling the extension of surrogate loss defined at Eq. (6) to matrix version $g_t(\cdot) : \mathcal{M}^{H+2} \mapsto \mathbb{R}$:

$$g_t(M) = \langle \nabla \tilde{f}(M_t), M \rangle - \mathbf{1}_{\{\langle \nabla \tilde{f}(M_t), V_t \rangle < 0\}} \cdot \langle \nabla \tilde{f}(M_t), V_t \rangle \cdot S_{\mathcal{M}}(M), \quad (24)$$

Algorithm 4 Efficient Control Algorithm for Dynamic Policy Regret

Input: Scream.Control algorithm \mathcal{A} .

- 1: Initialization: let \mathcal{A} project onto domain \mathcal{M}' ; submit $M_H \in \mathcal{M}$ in $[1, H]$ rounds.
 - 2: **for** $t = H + 1$ **to** T **do**
 - 3: Observe $c_t(\cdot, \cdot)$ and calculate the gradient of truncated loss function $\nabla \tilde{f}_t(M_t)$.
 - 4: Construct the surrogate loss $g_t(\cdot)$ according to Eq. (24).
 - 5: Compute the (sub-)gradient $\nabla g_t(M'_t)$ by Lemma 1 with extension to matrix.
 - 6: Send linearized loss $h_t(M) = \text{tr}(\nabla g_t(M'_t) \cdot M)$ to \mathcal{A} for update.
 - 7: Obtain decision M'_{t+1} from \mathcal{A} and submit $M_{t+1} = \Pi_{\mathcal{M}}[M'_{t+1}]$.
 - 8: **end for**
-

where $S_{\mathcal{M}}(M) = \inf_{A \in \mathcal{M}} \|A - M\|_F$ is the distance function to \mathcal{M} and we denote by $V_t = (M'_t - M_t) / \|M'_t - M_t\|_F$ the projection direction with M_t, M'_t defined in Algorithm 4. We denote by $\tilde{f}_t(\cdot)$ the unary truncated loss function constructed from cost $c_t(\cdot, \cdot)$ and refer the interested readers to Section 5.3 of Zhao et al. (2023) for more details, as in this paper we mainly focus on the projection issues. It is worth emphasizing that the truncated loss circumvents the growing of memory length with time, enabling the application of techniques from OCO with memory, while the gap between $\tilde{f}_t(\cdot)$ and $c_t(\cdot, \cdot)$ will not be too large.

The caveat to apply efficient method to Scream.Control remains that we should ensure the adoption of surrogate loss and surrogate domain will not ruin the transformation from non-stochastic control to OCO with memory, which requires the algorithm to account for the *switching cost* $\|M_{t-1} - M_t\|_F$ between parameters as well. Inspecting the algorithm derived from the efficient reduction closely, for two parameters M'_{t-1}, M'_t in the surrogate domain and the submitted parameters $M_{t-1} = \Pi_{\mathcal{M}}[M'_{t-1}], M_t = \Pi_{\mathcal{M}}[M'_t]$, the nonexpanding property of the projection operator in the Hilbert space implies that $\|\Pi_{\mathcal{M}}[M'_{t-1}] - \Pi_{\mathcal{M}}[M'_t]\|_F \leq \|M'_{t-1} - M'_t\|_F$ (Nemirovski et al., 2009), meaning that the switching cost $\|M_{t-1} - M_t\|_F$ can be controlled when applying efficient reduction mechanism. Therefore we can safely apply the efficient reduction to Scream.Control and improve the projection efficiency. Algorithm 4 enjoys the following theoretical guarantee with the proof sketch presented in Appendix E.1.

Theorem 6. *Under Assumptions 4-6, by choosing $H = \Theta(\log T)$, Algorithm 4 enjoys the dynamic regret of $\sum_{t=1}^T c_t(x_t, u_t) - \sum_{t=1}^T c_t(x_t^{\pi_t}, u_t^{\pi_t}) \leq \tilde{\mathcal{O}}(\sqrt{T(1 + P_T)})$, with one projection per round. Here, the comparators can be any feasible policies in $\Pi = \{\pi(K, M) \mid M \in \mathcal{M}\}$ with $\pi_t = \pi(K, M_t^*)$ for $t \in [T]$, and $P_T = \sum_{t=2}^T \|M_{t-1}^* - M_t^*\|_F$ measures the variation.*

Remark 5. The original algorithm of Zhao et al. (2023) requires maintaining $\mathcal{O}(\log T)$ base-learners, resulting in a computational cost of $\mathcal{O}(H \cdot \log T \cdot \min\{d_u^2 d_x, d_u d_x^2\})$ per round. In our variant algorithm, it permits the base-learners to project onto \mathcal{M}' , which can be achieved by simply rescaling the matrix. Consequently, the computational complexity can be significantly reduced by a factor of $\mathcal{O}(\log T)$ as depicted in Algorithm 4.

Remark 6. Note that this part focuses on optimizing dynamic regret of online non-stochastic control, rather than interval dynamic regret. We now explain why. Due to the switching cost, optimizing non-stationary regret is significantly more challenging than in standard OCO. The meta-base structure studied in this paper can be used to optimize the *dynamic regret* with injected correction terms, as demonstrated in (Zhao et al., 2023). On the other hand,

optimizing *adaptive regret* of OCO with switching cost requires a different approach. The current method (Zhang et al., 2022b) builds on parameter-free online learning (Cutkosky, 2020) to optimize the adaptive regret and tolerate the switching cost (so it may be possible to apply our reduction mechanism to (Cutkosky, 2020), as discussed in Appendix D.1, to reduce certain projection complexities). However, it remains unclear whether the meta-base approach can effectively optimize adaptive regret. Therefore, we focus on dynamic regret to streamline the projection complexity of the algorithm proposed in (Zhao et al., 2023). The result demonstrates the effectiveness and practicality of reducing projection complexity, especially given the complications of the feasible domains in non-stochastic control.

5.2 Online Principal Component Analysis

Principal Component Analysis (PCA) is a crucial dimensionality reduction technique, widely used in data processing, machine learning, and many more. Unlike the conventional offline PCA, online PCA is designed for scenarios where data arrive sequentially, thereby necessitating to conduct dimensionality reduction in an online manner.

5.2.1 PROBLEM FORMULATION

Online (uncentered-)PCA problem requires algorithms to forecast the optimal projection subspace upon receiving a series of streaming data on the fly (Warmuth and Kuzmin, 2008; Arora et al., 2013; Nie et al., 2016). Concretely, at each time t the algorithm receives an instance $\mathbf{x}_t \in \mathbb{R}^d$ (or in a more general setting, the algorithm receives an instance matrix $\mathbf{X}_t \in \mathbb{R}^{d \times d}$) and needs to project it onto a k -dimensional subspace ($k < d$) represented by a rank- k projection matrix $\mathbf{P}_t \in \mathcal{P}_k$. The domain of rank- k projection matrices is defined as

$$\mathcal{P}_k = \left\{ \mathbf{P} \in \mathbb{S}^d \mid \sigma_i(\mathbf{P}) \in \{0, 1\}, \text{rank}(\mathbf{P}) = k \right\}, \quad (25)$$

where \mathbb{S}^d denotes the set of real-valued $d \times d$ symmetric matrices and $\sigma_i(\cdot)$ denotes the i -th eigenvalue of the given matrix.

Online PCA uses *compression loss* $f_t(\mathbf{P}) = \|\mathbf{P}\mathbf{x}_t - \mathbf{x}_t\|_2^2$ to measure the reconstruction error at round t . Many prior works consider minimizing static regret for the online PCA problem, which benchmarks the cumulative compression loss of the learner against the fixed projection matrix in hindsight. However, in many environments the online data distributions can change over time, it is crucial to consider the non-stationarity issue in the algorithmic design. To this end, we investigate the adaptive regret for online PCA, which requires the algorithm to perform well for any interval $I \subseteq [T]$ with length $\tau = |I|$, defined as

$$\text{A-REG}(|I|) = \max_{[r, r+\tau-1] \subseteq [T]} \left\{ \sum_{t=r}^{r+\tau-1} f_t(\mathbf{P}_t) - \min_{\mathbf{P} \in \mathcal{P}_k} \sum_{t=r}^{r+\tau-1} f_t(\mathbf{P}) \right\}. \quad (26)$$

Yuan and Lamperski (2019) examine a regret notion similar to but weaker than Eq. (26), defined as $\text{WA-REG}(T) = \max_{[r, q] \subseteq [T]} \left\{ \sum_{t=r}^q f_t(\mathbf{P}_t) - \min_{\mathbf{P} \in \mathcal{P}_k} \sum_{t=r}^q f_t(\mathbf{P}) \right\}$. This variant is usually termed as the *weakly* adaptive regret (Hazan and Seshadhri, 2007; Zhang et al., 2018b), which lacks the guarantee for intervals with $|I| \leq \mathcal{O}(\sqrt{T})$. It should be noted that Yuan and Lamperski (2019) propose an algorithm with an $\tilde{\mathcal{O}}(\sqrt{T})$ weakly adaptive

regret for online PCA, and to the best of our knowledge, current literature lacks algorithms with an $\tilde{\mathcal{O}}(\sqrt{|I|})$ strongly adaptive regret for online PCA. We not only design the first algorithm with such a strongly adaptive regret guarantee, but also implement our reduction mechanism to ensure that it enjoys a projection complexity of 1.

5.2.2 PROJECTION COMPUTATIONAL COMPLEXITY

Before delving into the efficient projection mechanism concerned in our paper, we initiate with a brief overview on the projection challenge (due to the non-convexity issue) and its solution in the existing online PCA literature.

Notice that the feasible domain \mathcal{P}_k defined in Eq. (25) is inherently *non-convex*, making it hard to apply OCO techniques. To remedy this, the convex hull of \mathcal{P}_k is usually employed as a surrogate during the update, defined as $\hat{\mathcal{P}}_k = \{\mathbf{P} \in \mathbb{S}^d \mid \mathbf{0} \preceq \mathbf{P} \preceq \mathbf{I}, \text{tr}(\mathbf{P}) = k\}$. Nonetheless, the online PCA protocol requires the algorithm provide the decision within \mathcal{P}_k . To this end, the pioneering study of Warmuth and Kuzmin (2008) decomposes $\hat{\mathbf{P}}$ into a convex combination of, at most, d rank- k projection matrices represented as $\hat{\mathbf{P}} = \sum_{i=1}^d \lambda_i \mathbf{P}_i$, where $\lambda_i \in [0, 1]$ constitutes a distribution, and each $\mathbf{P}_i \in \mathcal{P}_k$ is a rank- k projection matrix. Following this decomposition, one can leverage the composite coefficients λ_i to sample a projection matrix as the submitted decision.

The gradient descent algorithm (Nie et al., 2016) can obtain $\mathcal{O}(\sqrt{kT})$ static regret for online PCA, which mainly consists of the following steps:

$$\hat{\mathbf{P}}'_{t+1} = \hat{\mathbf{P}}_t - \eta \nabla f_t(\mathbf{P}_t), \quad \hat{\mathbf{P}}_{t+1} = \arg \min_{\mathbf{P} \in \hat{\mathcal{P}}_k} \|\mathbf{P} - \hat{\mathbf{P}}'_{t+1}\|_F,$$

where $\eta > 0$ represents the step size and $\nabla f_t(\mathbf{P}_t)$ denotes the gradient with respect to \mathbf{P}_t . Then the algorithm samples a rank- k projection matrix based on $\hat{\mathbf{P}}_{t+1}$ to submit. The gradient descent step requires $\mathcal{O}(d^2)$ time expense, while the primary bottleneck is the projection step onto $\hat{\mathcal{P}}_k$, which typically demands $\mathcal{O}(d^3)$ computational complexity, owing to the matrix diagonalization process, as illustrated by Lemma 13 in Appendix E.2.

5.2.3 EFFICIENT REDUCTION

In this part, we provide our efficient algorithm for online PCA with the strongly adaptive regret guarantee. The algorithm is presented in Algorithm 5 and we introduce the necessary components in the below paragraphs.

By inspecting the convex surrogate domain $\hat{\mathcal{P}}_k$ carefully, onto which the algorithm projects, we propose the surrogate domain $\hat{\mathcal{P}}_k^s$ defined by the Frobenius norm as,

$$\hat{\mathcal{P}}_k^s = \left\{ \mathbf{P} \in \mathbb{S}^d \mid \|\mathbf{P}\|_F \leq \sqrt{k} \right\}, \quad (27)$$

which admits a fast projection by a simple rescaling.

At first glance, the compression loss $f_t(\mathbf{P}) = \|\mathbf{P}\mathbf{x}_t - \mathbf{x}_t\|_2^2$ seems to be quadratic, but indeed it is coordinate-wise linear with parameter \mathbf{P} as shown below,

$$f_t(\mathbf{P}) = \|\mathbf{P}\mathbf{x}_t - \mathbf{x}_t\|_2^2 = \text{tr}((\mathbf{I} - \mathbf{P})^2 \mathbf{x}_t \mathbf{x}_t^\top) = \text{tr}((\mathbf{I} - \mathbf{P}) \mathbf{x}_t \mathbf{x}_t^\top), \quad (28)$$

where the second equality is by the property of projection matrix \mathbf{P} . To ensure our algorithm's adaptability across varied scenarios, we consider a general setting, where the algorithm receives

Algorithm 5 Efficient Algorithm for Adaptive Regret under PCA Setting

-
- 1: Initialization: let $\mathbf{P}_1 = \hat{\mathbf{P}}_1$ be any point in \mathcal{P}_k ; let A_t denote the set of indexes for the active base-learners at time t .
 - 2: **for** $t = 1$ **to** T **do**
 - 3: Receive instance matrix $\mathbf{X}_t \in \mathbb{S}^d$.
 - 4: Construct the surrogate loss $g_t(\cdot)$ according to (29).
 - 5: Compute the (sub-)gradient $\nabla g_t(\hat{\mathbf{P}}_t^s)$ by Lemma 1 with extension to this problem.
 - 6: Remove base-learners whose deactivating time is t according to \mathcal{C} defined at Eq. (10).
 - 7: Initialize base-learner whose start time is $t + 1$ and set the learning rate for her $\eta_{t+1} = \frac{k(d-k)}{d|I_{t+1}|}$ where I_{t+1} is the active span according to \mathcal{C} defined at Eq. (10).
 - 8: Send $\nabla g_t(\hat{\mathbf{P}}_t^s)$ to each base-learner for update.
 - 9: For each $i \in A_t$, the base-learner updates the decision within $\hat{\mathcal{P}}_k^s$ defined at Eq. (27),

$$\hat{\mathbf{P}}_{t+1,i}^{s,\prime} = \hat{\mathbf{P}}_{t,i}^s - \eta_i \nabla g_t(\hat{\mathbf{P}}_t^s), \quad \hat{\mathbf{P}}_{t+1,i}^s = \hat{\mathbf{P}}_{t+1,i}^{s,\prime} \left(\mathbb{1}_{\{\|\hat{\mathbf{P}}_{t+1,i}^{s,\prime}\|_F \leq \sqrt{k}\}} + \frac{\sqrt{k}}{\|\hat{\mathbf{P}}_{t+1}^s\|_F} \mathbb{1}_{\{\|\hat{\mathbf{P}}_{t+1,i}^{s,\prime}\|_F > \sqrt{k}\}} \right).$$
 - 10: Meta-algorithm updates weight $\mathbf{p}_{t+1} \in \Delta_{|A_{t+1}|}$ according to Eq. (12), Eq. (13), and Eq. (14) with $\hat{\ell}_t = \text{tr}(\nabla g_t(\hat{\mathbf{P}}_t^s) \cdot \hat{\mathbf{P}}_t^s) / 2\sqrt{k}$ and $\ell_{t,i} = \text{tr}(\nabla g_t(\hat{\mathbf{P}}_t^s) \cdot \hat{\mathbf{P}}_{t,i}^s) / 2\sqrt{k}$.
 - 11: Compute $\hat{\mathbf{P}}_{t+1}^s = \sum_{i \in A_{t+1}} p_{t+1,i} \cdot \hat{\mathbf{P}}_{t+1,i}^s$, and $\hat{\mathbf{P}}_{t+1} = \Pi_{\hat{\mathcal{P}}_k}[\hat{\mathbf{P}}_{t+1}^s]$.
 - 12: Sample a projection matrix $\mathbf{P}_{t+1} \sim \hat{\mathbf{P}}_{t+1}$ to submit.
 - 13: **end for**
-

any semi-positive matrix $\mathbf{X}_t \in \mathbb{R}^{d \times d}$ as input instance rather than vector. Therefore, the loss function considered at Eq. (28) becomes $f_t(\mathbf{P}) = \text{tr}((\mathbf{I} - \mathbf{P})\mathbf{X}_t)$. We extend the surrogate loss defined at Eq. (6) to online PCA as

$$g_t(\mathbf{P}) = \text{tr}(\nabla f_t(\hat{\mathbf{P}}_t) \cdot \mathbf{P}) - \mathbb{1}_{\{\text{tr}(\nabla f_t(\hat{\mathbf{P}}_t) \cdot \mathbf{V}_t) < 0\}} \cdot \text{tr}(\nabla f_t(\hat{\mathbf{P}}_t) \cdot \mathbf{V}_t) \cdot S_{\hat{\mathcal{P}}_k}(\mathbf{P}), \quad (29)$$

where $S_{\hat{\mathcal{P}}_k}(\mathbf{P}) = \inf_{\mathbf{Q} \in \hat{\mathcal{P}}_k} \|\mathbf{P} - \mathbf{Q}\|_F$ is the minimum distance from \mathbf{P} to the domain $\hat{\mathcal{P}}_k$ and $\mathbf{V}_t = (\hat{\mathbf{P}}_t - \hat{\mathbf{P}}_t^s) / \|\hat{\mathbf{P}}_t - \hat{\mathbf{P}}_t^s\|_F$ denotes the matrix indicating the projection direction with $\hat{\mathbf{P}}_t, \hat{\mathbf{P}}_t^s$ defined in Algorithm 5. This surrogate loss enjoys the benign properties, as illustrated by Theorem 1 and Lemma 1. These two theoretical results are indeed consistent with the nearest-point projection in the Hilbert space. As for online PCA, the loss function is defined by the trace operator, and we employ the Frobenius norm as the distance metric for projection, which implies our optimization operates within a Hilbert space, making the aforementioned results directly applicable.

Given the absence of a strongly adaptive algorithm for the online PCA problem in the literature, we offer a detailed description of our algorithm (as opposed to the black-box reduction style for online non-stochastic control). We propose the algorithm by incorporating the gradient descent method (Nie et al., 2016) as the base-algorithm, Adapt-ML-Prod as the meta-algorithm, and the covering intervals defined at Eq. (10). Our efficient online PCA algorithm satisfies the following theorem, and we provide a proof sketch in Appendix E.2:

Theorem 7. *Assuming that $\|\mathbf{X}_t\|_F \leq 1$ for any $t \in [T]$ and $k \leq \frac{d}{2}$, then Algorithm 5 requires only one projection onto the domain $\hat{\mathcal{P}}_k$ per round and enjoys the following adaptive regret for any interval $I = [r, s] \subseteq [T]$: $\mathbb{E}[\sum_{t=r}^s f_t(\mathbf{P}_t)] - \min_{\mathbf{P} \in \mathcal{P}_k} \sum_{t=r}^s f_t(\mathbf{P}) \leq \tilde{\mathcal{O}}(\sqrt{k \cdot |I|})$, where we adopt the general setting by choosing $f_t(\mathbf{P}) = \text{tr}((\mathbf{I} - \mathbf{P})\mathbf{X}_t)$ and the expectation is due to the randomness introduced by the sampling of the algorithm.*

Remark 7. The projection operation onto $\hat{\mathcal{P}}_k$ is dominated by the matrix diagonalization which is of $\mathcal{O}(d^3)$ under general instances assumption. The vanilla adaptive PCA algorithm incurs $\mathcal{O}(d^3 \log T)$ computational cost by maintaining $\mathcal{O}(\log T)$ base-learners. Our efficient algorithm requires one projection and improves the computational cost to $\mathcal{O}(d^3)$ per round.

Remark 8. For online PCA, we focus on strongly adaptive regret to provide a strict improvement over Yuan and Lamperski (2019) — we achieve a reduction in projection complexity from $\mathcal{O}(\log T)$ to 1 per round and meanwhile improving the regret guarantee (previous result only considers weakly adaptive regret), significantly enhancing both computational efficiency and regret performance. While extending the analysis to *interval dynamic regret* for online PCA is feasible, it would complicate the process. Instead, we focus on adaptive regret to highlight the primary contributions of our method, particularly in terms of projection efficiency, which is crucial given the complexity of the feasible domain in online PCA.

6. Experiment

In this section, we provide empirical studies to evaluate our proposed methods.

General Setup. We conduct experiments on the synthetic data. We consider the following online regression problem. Let T denote the number of total rounds. At each round $t \in [T]$ the learner outputs the model parameter $\mathbf{w}_t \in \mathcal{W} \subseteq \mathbb{R}^d$ and simultaneously receives a data sample (x_t, y_t) with $x_t \in \mathcal{X} \subseteq \mathbb{R}^d$ being the feature and $y_t \in \mathbb{R}$ being the corresponding label.⁴ The learner can then evaluate her model by the online loss $f_t(\mathbf{w}_t) = \frac{1}{2}(x_t^\top \mathbf{w}_t - y_t)^2$ which uses a square loss to evaluate the difference between the predictive value $x_t^\top \mathbf{w}_t$ and the ground-truth label y_t , and then use the feedback information to update the model. In the simulations, we set $T = 20000$, the domain diameter as $D = 6$, and the dimension of the domain as $d = 8$. The feasible domain \mathcal{W} is set as an ellipsoid $\mathcal{W} = \{\mathbf{w} \in \mathbb{R}^d \mid \mathbf{w}^\top \mathbf{E} \mathbf{w} \leq \lambda_{\min}(\mathbf{E}) \cdot (D/2)^2\}$, where \mathbf{E} is a certain diagonal matrix and $\lambda_{\min}(\mathbf{E})$ denotes its minimum eigenvalue. Then, a projection onto \mathcal{W} requires solving a convex programming that is generally expensive. In the experiment, we use `scipy.optimize.NonlinearConstraint` to solve it to perform the projection onto the feasible domain.

To simulate the non-stationary online environments, we control the way to generate the data samples $\{(x_t, y_t)\}_{t=1}^T$. Specifically, for $t \in [T]$, the feature x_t is randomly sampled in an Euclidean ball with a diameter D same as the feasible domain of model parameters; and the corresponding label is set as $y_t = x_t^\top \mathbf{w}_t^* + \varepsilon_t$, where ε_t is the random noise drawn from $[0, 0.1]$ and \mathbf{w}_t^* is the underlying ground-truth model from the feasible domain \mathcal{W} generated according to a certain strategy specified below. For dynamic regret minimization, we simulate

4. With a slight abuse of notations, we here use \mathbf{w} to denote the model parameter and \mathcal{W} to denote the feasible domain, and meanwhile we reserve the notations of x and \mathcal{X} to denote the feature and feature space following the conventional notations of machine learning terminologies.

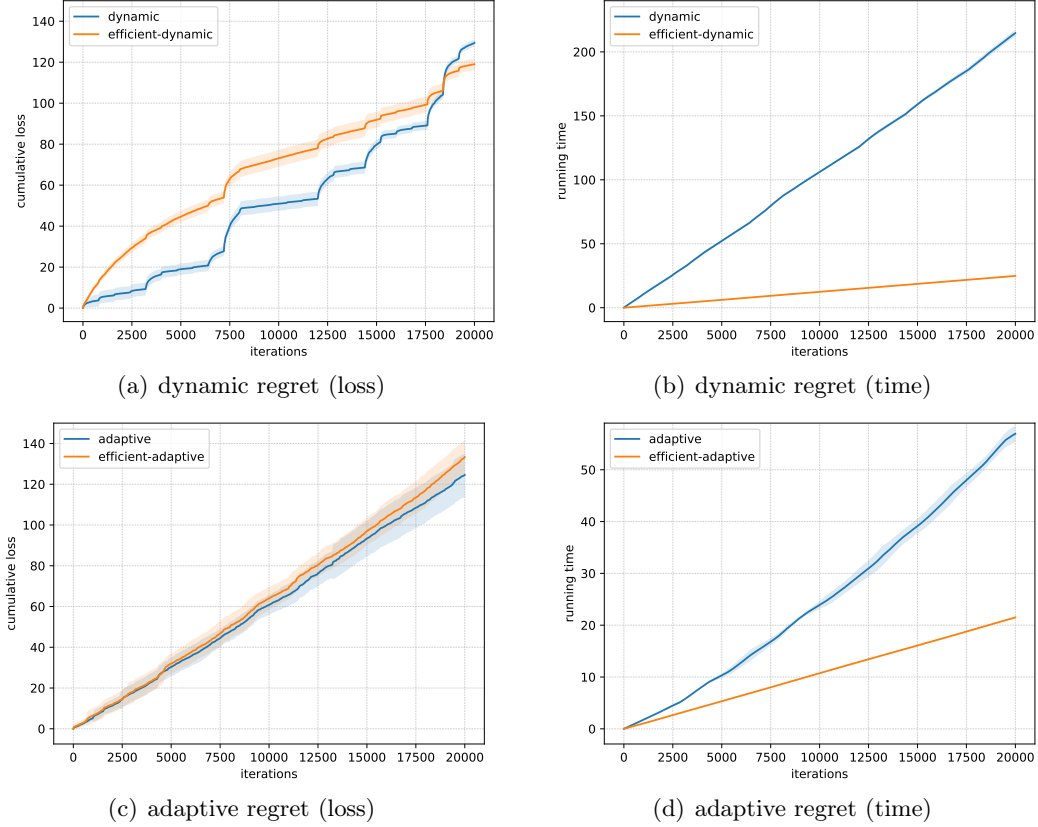


Figure 3: Performance comparisons of existing methods and our methods (indicated by “efficient” prefix) in terms of cumulative loss and running time (in seconds). The first two figures plot the results of methods for dynamic regret minimization, while the latter ones are for adaptive regret.

piecewise-stationary model drifts, as dynamic regret will be linear in T and thus vacuous when the model drift happens every round due to a linear path length measure. Concretely, we split the time horizon evenly into 25 stages and restrict the underlying model parameter \mathbf{w}_t^* to be stationary within a stage. For adaptive regret minimization, we simulate *gradually evolving* model drifts, where the underlying model parameter \mathbf{w}_{t+1}^* is generated based on the last-round model parameter \mathbf{w}_t^* with an additional random walk in the feasible domain \mathcal{W} . The step size of random walk is set proportional to D/T to ensure a smooth model change.

Contenders. For both dynamic regret and adaptive regret, we directly work on the small-loss online methods. We choose the Sword algorithm (Zhao et al., 2024) as the contender of our efficient method for dynamic regret (Algorithm 1) and choose the SACS algorithm (Zhang et al., 2019) as the contender of our efficient method for adaptive regret (Algorithm 2).

Results. We repeat the experiments for five times with different random seeds and report the results (mean and standard deviation) in Figure 3. We use a machine with a single CPU (Intel(R) Core(TM) i9-10900K CPU @ 3.70GHz) and 32GB main memory to conduct the experiments. We plot both cumulative loss and running time (in seconds) for all the

methods. We first examine the performance of dynamic regret minimization, see Figure 3(a) for cumulative loss and see Figure 3(b) for running time. The empirical results show that our method has a comparable performance to Sword without much sacrifice of cumulative loss, while achieving about 10 times speedup due to the improved projection complexity. Second, as shown in Figure 3(c) and Figure 3(d), a similar performance enhancement also appears in adaptive regret minimization, though the speedup is slightly smaller due to the fact that fewer learners are required to maintain for adaptive regret. To summarize, the empirical results show the effectiveness of our methods in retaining regret performance and also the efficiency in terms of running time due to the reduced projection complexity.

7. Conclusion

In this paper, we design a generic reduction mechanism that can reduce the projection complexity of two-layer methods for non-stationary online learning, thereby approaching a clearer resolution of necessary computational overhead for robustness to non-stationarity. Building on the reduction mechanism, we develop a collection of online algorithms that optimize dynamic regret, adaptive regret, and interval dynamic regret. All the algorithms retain the best known regret guarantees, and more importantly, require a single projection onto the feasible domain per iteration. Notably, due to the requirement of our reduction, all our algorithms only perform one gradient query and one function evaluation at each round as well, making them particularly attractive in scenarios with limited feedback and a need for lightweight updates. Furthermore, we present two applications with light project complexity, including online non-stochastic control and online principal component analysis. Finally, our empirical studies also corroborate the theoretical findings.

One important open question remains regarding another type of problem-dependent bound that scales with gradient variation, defined as $V_T \triangleq \sum_{t=2}^T \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x}) - \nabla f_{t-1}(\mathbf{x})\|_2^2$ (Chiang et al., 2012; Yang et al., 2014). This bound plays a crucial role in achieving fast convergence in zero-sum games (Syrkanis et al., 2015; Zhang et al., 2022a). While Zhao et al. (2024) introduced a two-layer method that attains gradient-variation dynamic regret with only one gradient and one function value query per iteration, developing a projection-efficient variant in this context remains an open challenge. As mentioned, the reduction requirement is necessary but not sufficient, and integrating optimistic online learning into our reduction mechanism is notably difficult due to the constrained feasible domain and the complexity of the two-layer structure. Another important problem is to understand the minimal computational overhead required for robustness to non-stationarity. For instance, is it possible to design a single-layer algorithm that achieves optimal dynamic regret, adaptive regret, or even interval dynamic regret, similar to the ensemble methods presented in this paper? Preliminary results suggest that a single-layer algorithm with a wavelet-based restarting scheme can achieve optimal dynamic regret in certain regimes, provided the online learner has access to noisy feedback about the comparators (Qian et al., 2024). However, its applicability to more general settings remains unclear. Extending this approach to broader scenarios would be valuable. Alternatively, information-theoretic arguments demonstrating the necessity of more complex structures would also be highly interesting.

Acknowledgment

This work was supported by National Science and Technology Major Project (2022ZD0114800) and NSFC (U23A20382). Peng Zhao was supported in part by the Xiaomi Foundation. We are grateful for anonymous reviewers for their careful reviews and many constructive suggestions.

References

- Dmitry Adamskiy, Wouter M. Koolen, Alexey V. Chernov, and Vladimir Vovk. A closer look at adaptive regret. *Journal of Machine Learning Research*, 17:23:1–23:21, 2016.
- Naman Agarwal, Brian Bullins, Elad Hazan, Sham M. Kakade, and Karan Singh. Online control with adversarial disturbances. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, pages 111–119, 2019a.
- Naman Agarwal, Elad Hazan, and Karan Singh. Logarithmic regret for online control. In *Advances in Neural Information Processing Systems 32 (NeurIPS)*, pages 10175–10184, 2019b.
- Oren Anava, Elad Hazan, and Shie Mannor. Online learning for adversaries with memory: Price of past mistakes. In *Advances in Neural Information Processing Systems 28 (NIPS)*, pages 784–792, 2015.
- Raman Arora, Andrew Cotter, and Nati Srebro. Stochastic optimization of PCA with capped MSG. In *Advances in Neural Information Processing Systems 26 (NIPS)*, pages 1815–1823, 2013.
- Peter Auer, Nicolò Cesa-Bianchi, and Claudio Gentile. Adaptive and self-confident on-line learning algorithms. *Journal of Computer and System Sciences*, 64(1):48–75, 2002.
- Dheeraj Baby and Yu-Xiang Wang. Optimal dynamic regret in exp-concave online learning. In *Proceedings of the 34th Conference on Learning Theory (COLT)*, pages 359–409, 2021.
- Yong Bai, Yu-Jie Zhang, Peng Zhao, Masashi Sugiyama, and Zhi-Hua Zhou. Adapting to online label shift with provable guarantees. In *Advances in Neural Information Processing Systems 35 (NeurIPS)*, pages 29960–29974, 2022.
- Chao-Kai Chiang, Tianbao Yang, Chia-Jung Lee, Mehrdad Mahdavi, Chi-Jen Lu, Rong Jin, and Shenghuo Zhu. Online optimization with gradual variations. In *Proceedings of the 25th Conference On Learning Theory (COLT)*, pages 6.1–6.20, 2012.
- Ashok Cutkosky. Parameter-free, dynamic, and strongly-adaptive online learning. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, pages 2250–2259, 2020.
- Ashok Cutkosky and Francesco Orabona. Black-box reductions for parameter-free online learning in banach spaces. In *Proceedings of the 31st Conference on Learning Theory (COLT)*, pages 1493–1529, 2018.

- Amit Daniely, Alon Gonen, and Shai Shalev-Shwartz. Strongly adaptive online learning. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, pages 1405–1411, 2015.
- Dylan J. Foster and Max Simchowitz. Logarithmic regret for adversarial online control. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, pages 3211–3221, 2020.
- Yoav Freund and Robert E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997.
- Pierre Gaillard, Gilles Stoltz, and Tim van Erven. A second-order bound with excess losses. In *Proceedings of the 27th Conference on Learning Theory (COLT)*, pages 176–196, 2014.
- Paula Gradu, John Hallman, and Elad Hazan. Non-stochastic control with bandit feedback. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, pages 10764–10774, 2020.
- András György and Csaba Szepesvári. Shifting regret, mirror descent, and matrices. In *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, pages 2943–2951, 2016.
- Eric C. Hall and Rebecca Willett. Dynamical models and tracking regret in online convex programming. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, pages 579–587, 2013.
- Elad Hazan. Introduction to Online Convex Optimization. *Foundations and Trends in Optimization*, 2(3-4):157–325, 2016.
- Elad Hazan and C. Seshadhri. Adaptive algorithms for online decision problems. *Electronic Colloquium on Computational Complexity (ECCC)*, 14(088), 2007.
- Elad Hazan and C. Seshadhri. Efficient learning algorithms for changing environments. In *Proceedings of the 26th International Conference on Machine Learning (ICML)*, pages 393–400, 2009.
- Elad Hazan and Karan Singh. Introduction to online nonstochastic control. *ArXiv preprint*, arXiv:2211.09619, 2022.
- Elad Hazan, Sham M. Kakade, and Karan Singh. The nonstochastic control problem. In *Proceedings of the 31st International Conference on Algorithmic Learning Theory (ALT)*, pages 408–421, 2020.
- Andrew Jacobsen and Ashok Cutkosky. Parameter-free mirror descent. In *Proceedings of the 35th Conference on Learning Theory (COLT)*, pages 4160–4211, 2022.
- Kwang-Sung Jun, Francesco Orabona, Stephen Wright, and Rebecca Willett. Improved strongly adaptive online learning using coin betting. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 943–951, 2017.

- Rudolf Emil Kalman. Contributions to the theory of optimal control. *Bol. Soc. Mat. Mexicana*, 5(2):102–119, 1960.
- Long-Fei Li, Peng Zhao, and Zhi-Hua Zhou. Near-optimal dynamic regret for adversarial linear mixture mdps. In *Advances in Neural Information Processing Systems 37 (NeurIPS)*, pages 55858–55883, 2024.
- Haipeng Luo and Robert E. Schapire. Achieving all with no parameters: AdaNormalHedge. In *Proceedings of the 28th Annual Conference Computational Learning Theory (COLT)*, pages 1286–1304, 2015.
- Zakaria Mhammedi, Wouter M. Koolen, and Tim van Erven. Lipschitz adaptivity with multiple learning rates in online learning. In *Proceedings of the 32nd Annual Conference Computational Learning Theory (COLT)*, pages 2490–2511, 2019.
- Arkadi Nemirovski, Anatoli Juditsky, Guanghui Lan, and Alexander Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.
- Jiazhong Nie, Wojciech Kotłowski, and Manfred K. Warmuth. Online PCA with optimal regret. *Journal of Machine Learning Research*, 17:173:1–173:49, 2016.
- Francesco Orabona. A Modern Introduction to Online Learning. *ArXiv preprint*, 1912.13213, 2019.
- Francesco Orabona and Dávid Pál. Scale-free online learning. *Theoretical Computer Science*, 716:50–69, 2018.
- Francesco Orabona, Nicolò Cesa-Bianchi, and Claudio Gentile. Beyond logarithmic bounds in online learning. In *Proceedings of the 15th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 823–831, 2012.
- Roman Pogodin and Tor Lattimore. On first-order bounds, variance and gap-dependent bounds for adversarial bandits. In *Proceedings of the 35th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 894–904, 2019.
- Yu-Yang Qian, Peng Zhao, Yu-Jie Zhang Masashi Sugiyama, and Zhi-Hua Zhou. Efficient non-stationary online learning by wavelets with applications to online distribution shift adaptation. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, pages 41383–41415, 2024.
- Shai Shalev-Shwartz. Online Learning: Theory, Algorithms and Applications. *PhD Thesis*, 2007.
- Shai Shalev-Shwartz. Online Learning and Online Convex Optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2012.
- Nathan Srebro, Karthik Sridharan, and Ambuj Tewari. Smoothness, low noise and fast rates. In *Advances in Neural Information Processing Systems 23 (NIPS)*, pages 2199–2207, 2010.

- Vasilis Syrgkanis, Alekh Agarwal, Haipeng Luo, and Robert E. Schapire. Fast convergence of regularized learning in games. In *Advances in Neural Information Processing Systems 28 (NIPS)*, pages 2989–2997, 2015.
- Tim van Erven and Wouter M. Koolen. Metagrad: Multiple learning rates in online learning. In *Advances in Neural Information Processing Systems 29 (NIPS)*, pages 3666–3674, 2016.
- Guanghai Wang, Dakuan Zhao, and Lijun Zhang. Minimizing adaptive regret with one gradient per iteration. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2762–2768, 2018.
- Guanghai Wang, Shiyin Lu, and Lijun Zhang. Adaptivity and optimality: A universal algorithm for online convex optimization. In *Proceedings of the 35th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 659–668, 2019.
- Manfred K Warmuth and Dima Kuzmin. Randomized online pca algorithms with regret bounds that are logarithmic in the dimension. *Journal of Machine Learning Research*, 9 (Oct):2287–2320, 2008.
- Tianbao Yang, Mehrdad Mahdavi, Rong Jin, and Shenghuo Zhu. Regret bounded by gradual variation for online convex optimization. *Machine Learning*, 95(2):183–223, 2014.
- Wenhao Yang, Wei Jiang, Yibo Wang, Ping Yang, Yao Hu, and Lijun Zhang. Small-loss adaptive regret for online convex optimization. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, pages 56156–56195, 2024.
- Jianjun Yuan and Andrew G. Lamperski. Online adaptive principal component analysis and its extensions. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning (ICML)*, volume 97, pages 7213–7221, 2019.
- Lijun Zhang. Online learning in changing environments. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 5178–5182, 2020. Early Career.
- Lijun Zhang, Shiyin Lu, and Zhi-Hua Zhou. Adaptive online learning in dynamic environments. In *Advances in Neural Information Processing Systems 31 (NeurIPS)*, pages 1330–1340, 2018a.
- Lijun Zhang, Tianbao Yang, Rong Jin, and Zhi-Hua Zhou. Dynamic regret of strongly adaptive methods. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, pages 5877–5886, 2018b.
- Lijun Zhang, Tie-Yan Liu, and Zhi-Hua Zhou. Adaptive regret of convex and smooth functions. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, pages 7414–7423, 2019.
- Lijun Zhang, Shiyin Lu, and Tianbao Yang. Minimizing dynamic regret and adaptive regret simultaneously. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 309–319, 2020.

- Lijun Zhang, Guanghui Wang, Wei-Wei Tu, and Zhi-Hua Zhou. Dual adaptivity: A universal algorithm for minimizing the adaptive regret of convex functions. In *Advances in Neural Information Processing Systems 34 (NeurIPS)*, pages 24968–24980, 2021.
- Mengxiao Zhang, Peng Zhao, Haipeng Luo, and Zhi-Hua Zhou. No-regret learning in time-varying zero-sum games. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*, pages 26772–26808, 2022a.
- Yu-Jie Zhang, Zhen-Yu Zhang, Peng Zhao, and Masashi Sugiyama. Adapting to continuous covariate shift via online density ratio estimation. In *Advances in Neural Information Processing Systems 36 (NeurIPS)*, pages 29074–29113, 2023.
- Yu-Jie Zhang, Peng Zhao, and Masashi Sugiyama. Non-stationary online learning for curved losses: Improved dynamic regret via mixability. In *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, page to appear, 2025.
- Zhiyu Zhang, Ashok Cutkosky, and Yannis Paschalidis. Optimal comparator adaptive online learning with switching cost. In *Advances in Neural Information Processing Systems 35 (NeurIPS)*, pages 23936–23950, 2022b.
- Peng Zhao, Yu-Jie Zhang, Lijun Zhang, and Zhi-Hua Zhou. Dynamic regret of convex and smooth functions. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, pages 12510–12520, 2020.
- Peng Zhao, Guanghui Wang, Lijun Zhang, and Zhi-Hua Zhou. Bandit convex optimization in non-stationary environments. *Journal of Machine Learning Research*, 22(125):1–45, 2021.
- Peng Zhao, Long-Fei Li, and Zhi-Hua Zhou. Dynamic regret of online markov decision processes. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*, pages 26865–26894, 2022a.
- Peng Zhao, Yu-Xiang Wang, and Zhi-Hua Zhou. Non-stationary online learning with memory and non-stochastic control. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 2101–2133, 2022b.
- Peng Zhao, Yan-Feng Xie, Lijun Zhang, and Zhi-Hua Zhou. Efficient methods for non-stationary online learning. In *Advances in Neural Information Processing Systems 35 (NeurIPS)*, pages 11573–11585, 2022c.
- Peng Zhao, Yu-Hu Yan, Yu-Xiang Wang, and Zhi-Hua Zhou. Non-stationary online learning with memory and non-stochastic control. *Journal of Machine Learning Research*, 24(206):1–70, 2023.
- Peng Zhao, Yu-Jie Zhang, Lijun Zhang, and Zhi-Hua Zhou. Adaptivity and non-stationarity: Problem-dependent dynamic regret for online convex optimization. *Journal of Machine Learning Research*, 25(98):1–52, 2024.
- Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning (ICML)*, pages 928–936, 2003.

Appendix A. Omitted Details for Reduction Mechanism

In this section, we provide the proofs of Theorem 1 and Lemma 1.

A.1 Properties of Distance Function

Before presenting the proofs, we here collect two useful lemmas regarding the distance function used in the surrogate loss, which will be useful in the following proofs. The proofs of the two lemmas can be found in the seminal paper of Cutkosky and Orabona (2018).

Lemma 2 (Proposition 1 of Cutkosky and Orabona (2018)). *The distance function $S_{\mathcal{X}}(\mathbf{y}) = \inf_{\mathbf{x} \in \mathcal{X}} \|\mathbf{y} - \mathbf{x}\|_2$ is convex and 1-Lipschitz for any closed convex feasible domain $\mathcal{X} \subseteq \mathbb{R}^d$.*

Lemma 3 (Theorem 4 of Cutkosky and Orabona (2018)). *Let $\mathcal{X} \subseteq \mathbb{R}^d$ be a closed convex set. Given $\mathbf{y} \in \mathbb{R}^d$ and $\mathbf{y} \notin \mathcal{X}$. Let $\mathbf{x} = \Pi_{\mathcal{X}}[\mathbf{y}]$. Then we have $\{\frac{\mathbf{y}-\mathbf{x}}{\|\mathbf{y}-\mathbf{x}\|_2}\} = \partial S_{\mathcal{X}}(\mathbf{y})$.*

A.2 Proof of Theorem 1

Theorem 1 is originally due to Cutkosky (2020), and for self-containedness we restate their proof using our notations.

Proof. When $\langle \nabla f_t(\mathbf{x}_t), \mathbf{v}_t \rangle \geq 0$, by the definition of the surrogate loss defined in Eq. (6), we have $g_t(\mathbf{y}) = \langle \nabla f_t(\mathbf{x}_t), \mathbf{y} \rangle$, which is linear in \mathbf{y} and thus convex. It is clear that $\|\nabla g_t(\mathbf{y}_t)\|_2 = \|\nabla f_t(\mathbf{x}_t)\|_2$ and thus satisfies the claimed inequality of gradient norms in the statement. Moreover, the inequality (7) holds evidently due to the linear surrogate loss in this case.

Let us focus on the case when $\langle \nabla f_t(\mathbf{x}_t), \mathbf{v}_t \rangle < 0$. First, it can be verified that the surrogate loss $g_t(\mathbf{y}) = \langle \nabla f_t(\mathbf{x}_t), \mathbf{y} \rangle - \langle \nabla f_t(\mathbf{x}_t), \mathbf{v}_t \rangle \cdot S_{\mathcal{X}}(\mathbf{y})$ is convex due to the convexity of $S_{\mathcal{X}}(\mathbf{y})$ shown in Lemma 2 and the condition of $\langle \nabla f_t(\mathbf{x}_t), \mathbf{v}_t \rangle < 0$ in this case. Next, the gradient of $g_t(\cdot)$ at the \mathbf{y}_t point can be calculated according to Lemma 1 as

$$\nabla g_t(\mathbf{y}_t) = \nabla f_t(\mathbf{x}_t) - \langle \nabla f_t(\mathbf{x}_t), \mathbf{v}_t \rangle \cdot \mathbf{v}_t$$

where $\mathbf{v}_t = (\mathbf{y}_t - \mathbf{x}_t) / \|\mathbf{y}_t - \mathbf{x}_t\|_2$. Notice that $\|\mathbf{v}_t\|_2 = 1$ and $\nabla g_t(\mathbf{y}_t)$ is an orthogonal projection of $\nabla f_t(\mathbf{x}_t)$ onto the subspace perpendicular to the vector \mathbf{v}_t , so we have $\|\nabla g_t(\mathbf{y}_t)\|_2 \leq \|\nabla f_t(\mathbf{x}_t)\|_2$. Finally, we proceed to prove the inequality (7) in this case. Since the comparator $\mathbf{u}_t \in \mathcal{X}$ is in the feasible domain, we have $S_{\mathcal{X}}(\mathbf{u}_t) = \|\mathbf{u}_t - \mathbf{u}_t\|_2 = 0$ and get

$$\begin{aligned} & \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u}_t \rangle \\ &= \langle \nabla f_t(\mathbf{x}_t), \mathbf{y}_t \rangle + \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{y}_t \rangle - \langle \nabla f_t(\mathbf{x}_t), \mathbf{u}_t \rangle \\ &= \langle \nabla f_t(\mathbf{x}_t), \mathbf{y}_t \rangle - \langle \nabla f_t(\mathbf{x}_t), \frac{\mathbf{y}_t - \mathbf{x}_t}{\|\mathbf{y}_t - \mathbf{x}_t\|_2} \rangle \cdot \|\mathbf{y}_t - \mathbf{x}_t\|_2 - \langle \nabla f_t(\mathbf{x}_t), \mathbf{u}_t \rangle \\ &= \langle \nabla f_t(\mathbf{x}_t), \mathbf{y}_t \rangle - \langle \nabla f_t(\mathbf{x}_t), \mathbf{v}_t \rangle \cdot S_{\mathcal{X}}(\mathbf{y}_t) - \langle \nabla f_t(\mathbf{x}_t), \mathbf{u}_t \rangle + \langle \nabla f_t(\mathbf{x}_t), \mathbf{v}_t \rangle \cdot S_{\mathcal{X}}(\mathbf{u}_t) \\ &= g_t(\mathbf{y}_t) - g_t(\mathbf{u}_t) \\ &\leq \langle \nabla g_t(\mathbf{y}_t), \mathbf{y}_t - \mathbf{u}_t \rangle, \end{aligned}$$

where the last inequality holds owing to the convexity of the surrogate loss proven earlier.

Combining the two cases finishes the proof. \square

A.3 Proof of Lemma 1

Lemma 1 is originally due to Cutkosky and Orabona (2018), and for self-containedness we restate their proof using our notations.

Proof. With a slight abuse of notations, we simply use $\nabla g_t(\mathbf{y})$ to denote the (sub-)gradient of surrogate function $g_t(\cdot)$ at point \mathbf{y} , no matter whether the function is differentiable.

When $\langle \nabla f_t(\mathbf{x}_t), \mathbf{v}_t \rangle \geq 0$, the surrogate loss is $g_t(\mathbf{y}) = \langle \nabla f_t(\mathbf{x}_t), \mathbf{y} \rangle$ by definition in Eq. (6). Therefore, the gradient simply becomes $\nabla g_t(\mathbf{y}_t) = \nabla f_t(\mathbf{x}_t)$.

When $\langle \nabla f_t(\mathbf{x}_t), \mathbf{v}_t \rangle < 0$, the surrogate loss becomes $g_t(\mathbf{y}) = \langle \nabla f_t(\mathbf{x}_t), \mathbf{y} \rangle - \langle \nabla f_t(\mathbf{x}_t), \mathbf{v}_t \rangle \cdot S_{\mathcal{X}}(\mathbf{y})$ according to definition in Eq. (6). By Lemma 3, the gradient can be calculated by

$$\nabla g_t(\mathbf{y}) = \nabla f_t(\mathbf{x}_t) - \langle \nabla f_t(\mathbf{x}_t), \mathbf{v}_t \rangle \cdot \frac{\mathbf{y} - \Pi_{\mathcal{X}}[\mathbf{y}]}{\|\mathbf{y} - \Pi_{\mathcal{X}}[\mathbf{y}]\|_2},$$

where the computation needs the projection onto domain \mathcal{X} . In particular, for \mathbf{y}_t , we have

$$\nabla g_t(\mathbf{y}_t) = \nabla f_t(\mathbf{x}_t) - \langle \nabla f_t(\mathbf{x}_t), \mathbf{v}_t \rangle \cdot \frac{\mathbf{y}_t - \mathbf{x}_t}{\|\mathbf{y}_t - \mathbf{x}_t\|_2} = \nabla f_t(\mathbf{x}_t) - \langle \nabla f_t(\mathbf{x}_t), \mathbf{v}_t \rangle \cdot \mathbf{v}_t.$$

This ends the proof. \square

Appendix B. Omitted Details for Dynamic Regret Minimization

In this section, we provide the proofs for the theorems presented in Section 2. Specifically, we first prove the worst-case bound (Theorem 2) and then the small-loss bound (Theorem 3).

B.1 Proof of Theorem 2

Proof. Notice that Zhang et al. (2018a) propose the improved Ader algorithm (see Algorithm 3 and Algorithm 4 in their paper), which uses the linearized loss as the input which allows the online algorithm to require one gradient and one function value. So the algorithm satisfies the requirements of our reduction mechanism, and our algorithm (Algorithm 1 with specifications in Theorem 1) can be regarded as the improved Ader equipped with the projection-efficient reduction. As a consequence, we can directly obtain the same dynamic regret guarantee and meanwhile ensure 1 projection complexity by following the same proof of the improved Ader as well as the reduction guarantee (Theorem 1). \square

B.2 Proof of Theorem 3

Proof. By the reduction guarantee shown in Theorem 1, we have the following result that decomposes the dynamic regret into the two terms.

$$\begin{aligned} \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}_t) &\leq \sum_{t=1}^T g_t(\mathbf{y}_t) - \sum_{t=1}^T g_t(\mathbf{u}_t) \leq \sum_{t=1}^T \langle \nabla g_t(\mathbf{y}_t), \mathbf{y}_t - \mathbf{u}_t \rangle \\ &= \underbrace{\sum_{t=1}^T \langle \nabla g_t(\mathbf{y}_t), \mathbf{y}_t - \mathbf{y}_{t,i} \rangle}_{\text{meta-regret}} + \underbrace{\sum_{t=1}^T \langle \nabla g_t(\mathbf{y}_t), \mathbf{y}_{t,i} - \mathbf{u}_t \rangle}_{\text{base-regret}}, \end{aligned} \quad (30)$$

where in (30) the first term is called *meta-regret* as it measures the regret overhead of the meta-algorithm to track the unknown best base-learner, and the second term is called the *base-regret* to denote the dynamic regret of the base-learner i . Note that the above decomposition holds for any base-learner index $i \in [N]$.

Upper bound of meta-regret. Since the meta-algorithm is essentially FTRL with time-varying learning rates and a negative entropy regularizer, we apply Lemma 18 to obtain the meta-regret upper bound by choosing $\ell_{t,i} = \langle \nabla g_t(\mathbf{y}_t), \mathbf{y}_{t,i} \rangle$ and obtain that

$$\begin{aligned}
 \sum_{t=1}^T \langle \nabla g_t(\mathbf{y}_t), \mathbf{y}_t - \mathbf{y}_{t,i} \rangle &\leq 3 \sqrt{\ln N \left(1 + \sum_{t=1}^T D^2 \|\nabla g_t(\mathbf{y}_t)\|_2^2 \right)} + \frac{G^2 D^2 \sqrt{\ln N}}{2} \\
 &\leq 3D \sqrt{\ln N \sum_{t=1}^T \|\nabla g_t(\mathbf{y}_t)\|_2^2} + \frac{(6 + G^2 D^2) \sqrt{\ln N}}{2} \\
 &\leq 3D \sqrt{\ln N \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t)\|_2^2} + \mathcal{O}(1) \\
 &\leq 6D \sqrt{L \ln N \sum_{t=1}^T f_t(\mathbf{x}_t)} + \mathcal{O}(1), \tag{31}
 \end{aligned}$$

where the first inequality holds because we have $\|\ell_t\|_\infty^2 = \max_{i \in [N]} (\langle \nabla g_t(\mathbf{y}_t), \mathbf{y}_{t,i} \rangle)^2 \leq D^2 \|\nabla g_t(\mathbf{y}_t)\|_2^2$ by Cauchy-Schwarz inequality and $\|\nabla g_t(\mathbf{y}_t)\|_2 \leq \|\nabla f_t(\mathbf{x}_t)\|_2 \leq G$ by Theorem 1, and the last inequality is due to the self-bounding properties of smooth functions (see Lemma 20). Note that $\mathcal{O}(\ln N) = \mathcal{O}(\log \log T)$ can be treated as a constant following previous studies (Luo and Schapire, 2015; Gaillard et al., 2014; Zhao et al., 2024).

Upper bound of base-regret. According to Lemma 14 and noticing that the comparator sequence $\mathbf{u}_1, \dots, \mathbf{u}_T \in \mathcal{X} \subseteq \mathcal{Y}$ and the diameter of \mathcal{Y} equals to $2D$ by definition, with slight modifications, we have the following dynamic regret bound.

$$\begin{aligned}
 \sum_{t=1}^T \langle \nabla g_t(\mathbf{y}_t), \mathbf{y}_{t,i} - \mathbf{u}_t \rangle &\leq \frac{5D^2}{2\eta_i} + \frac{D}{\eta_i} \sum_{t=2}^T \|\mathbf{u}_t - \mathbf{u}_{t-1}\|_2 + \eta_i \sum_{t=1}^T \|\nabla g_t(\mathbf{y}_t)\|_2^2 \\
 &\leq \frac{5D^2}{2\eta_i} + \frac{D}{\eta_i} \sum_{t=2}^T \|\mathbf{u}_t - \mathbf{u}_{t-1}\|_2 + 4\eta_i L \sum_{t=1}^T f_t(\mathbf{x}_t),
 \end{aligned}$$

where the last inequality uses the property of $\|\nabla g_t(\mathbf{y}_t)\|_2 \leq \|\nabla f_t(\mathbf{x}_t)\|_2$ in Theorem 1 and the self-bounding property of smooth functions (see Lemma 20).

Upper bound of dynamic regret. Combining the upper bounds meta-regret and base-regret together yields the following dynamic regret:

$$\sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}_t) \leq 6D \sqrt{L \ln N \sum_{t=1}^T f_t(\mathbf{x}_t)} + \frac{5D^2 + 2DP_T}{2\eta_i} + 4\eta_i L \sum_{t=1}^T f_t(\mathbf{x}_t) + \mathcal{O}(1), \tag{32}$$

which holds for any base-learner's index $i \in [N]$.

Next, we specify the base-learner \mathcal{E}_i to be compared with. Indeed, we aim at choosing the one with step size closest to the (near-)optimal step size $\eta^* = \sqrt{\frac{5D^2 + 2DP_T}{1 + 8LF_T^x}}$, where we denote by $F_T^x = \sum_{t=1}^T f_t(\mathbf{x}_t)$ the cumulative loss of the decisions. By Assumption 1 and Assumption 2, we have $F_T^x \in [0, GDT]$ and then the possible minimum optimal and maximum step size are

$$\eta_{\min} = \sqrt{\frac{5D^2}{1 + 8LGD T}}, \text{ and } \eta_{\max} = \sqrt{5D^2 + 2D^2 T}.$$

The construction of step size pool is by discretizing the interval $[\eta_{\min}, \eta_{\max}]$ with intervals with exponentially increasing length. The step size of each base-learner is designed to be monotonically increasing with respect to the index. Consequently, it is evident to verify that there exists an index $i^* \in [N]$ such that $\eta_{i^*} \leq \eta^* \leq \eta_{i^*+1} = 2\eta_{i^*}$. As the upper bounds of meta-regret and base-regret hold for any compared base-learner, we can choose the index as i^* in particular. Then the second and the third terms in the inequality (32) satisfy

$$\begin{aligned} & \frac{5D^2 + 2DP_T}{2\eta_{i^*}} + 4\eta_{i^*}LF_T^x \\ & \leq \frac{5D^2 + 2DP_T}{\eta^*} + 4\eta^*LF_T^x \\ & \leq \sqrt{(5D^2 + 2DP_T)(1 + 8LF_T^x)} + \frac{1}{2}\sqrt{(5D^2 + 2DP_T)(1 + 8LF_T^x)} \\ & \leq 3\sqrt{2(5D^2 + 2DP_T)(1 + LF_T^x)}. \end{aligned} \tag{33}$$

Substituting inequality (33) into inequality (32), we have,

$$\begin{aligned} & \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}_t) \\ & \leq 6D\sqrt{L \ln NF_T^x} + 3\sqrt{2(5D^2 + 2DP_T)(1 + LF_T^x)} + \mathcal{O}(1) \\ & \leq \left(6D\sqrt{L \ln N} + 3\sqrt{2L(5D^2 + 2DP_T)}\right) \sqrt{F_T^x} + 3\sqrt{2(5D^2 + 2DP_T)} + \mathcal{O}(1) \\ & \leq \mathcal{O}\left(\sqrt{(1 + P_T)(F_T^u + \sqrt{P_T} + \mathcal{O}(1))} + P_T + 1\right) \\ & = \mathcal{O}\left(\sqrt{(F_T^u + P_T)(1 + P_T)}\right), \end{aligned}$$

where the last inequality holds by Lemma 25. Hence, we complete the proof of Theorem 3. \square

Appendix C. Omitted Details for Adaptive Regret Minimization

In this section, we present omitted details for minimizing the worst-case and small-loss adaptive regret bounds with a focus on proving the main theorem of small-loss bound, i.e., Theorem 4. Appendix C.1 provides key lemmas, Appendix C.2 presents the proof of Theorem 4, and Appendix C.3 – Appendix C.6 give the proofs of these key lemmas.

C.1 Key Lemmas

In this part, we present three key lemmas for proving Theorem 4, based on which we prove Theorem 4 in Appendix C.2. We will prove those lemmas in the following several subsections.

The first lemma gives the second-order regret bound for the meta-algorithm (Adapt-ML-Prod) (Gaillard et al., 2014), which plays a crucial role in applying our reduction. Though Adapt-ML-Prod can be applied to the sleeping-expert setting directly, we need more careful analysis to obtain the *fully* small-loss adaptive regret bound, otherwise the direct reduction of results from Gaillard et al. (2014) will incur an undesired $\mathcal{O}(\log T)$ factor.

Lemma 4. *Under Assumptions 2 and 3, for any interval $I = [i, j] \in \tilde{\mathcal{C}}$ in the geometric covers defined in Eq. (11) at the beginning of which we suppose the m -th base-learner is initialized, Algorithm 2 ensures*

$$\begin{aligned} \sum_{\tau=i}^t g_{\tau}(\mathbf{y}_{\tau}) - g_{\tau}(\mathbf{y}_{\tau,m}) &\leq \sum_{\tau=i}^t \langle \nabla g_{\tau}(\mathbf{y}_{\tau}), \mathbf{y}_{\tau} - \mathbf{y}_{\tau,m} \rangle \\ &\leq 4D \left(3\sqrt{\ln(1+2m)} + \frac{\mu_t}{\sqrt{\ln(1+2m)}} \right) \sqrt{L \sum_{\tau=i}^t f_{\tau}(\mathbf{x}_{\tau}) + 18GD \ln(1+2m) + 6GD\mu_t} \\ &= \mathcal{O} \left(\sqrt{\log(m) \sum_{\tau=i}^t f_{\tau}(\mathbf{x}_{\tau}) + \log(m)} \right), \end{aligned}$$

where we denote $\mu_t = \ln(1 + (1 + \ln(1 + t))/(2e))$.

Combining the above lemma with the regret bound for base-learners, we can obtain the adaptive regret for any interval in the geometric covering intervals $\tilde{\mathcal{C}}$ defined in Eq. (11).

Lemma 5. *Under Assumptions 1, 2, and 3, for any interval $[i, j] \in \tilde{\mathcal{C}}$ in the geometric covering intervals defined in Eq. (11), at the beginning of which we assume the m -th base-learner is initialized, Algorithm 2 ensures for any time $t \in [i, j]$ and any comparator $\mathbf{u} \in \mathcal{X}$,*

$$\begin{aligned} \sum_{\tau=i}^t f_{\tau}(\mathbf{x}_{\tau}) - \sum_{\tau=i}^t f_{\tau}(\mathbf{u}) &\leq 4D \left(3\sqrt{\ln(1+2m)} + \mu_t + 2 \right) \sqrt{L \sum_{\tau=i}^t f_{\tau}(\mathbf{x}_{\tau}) + 18GD \ln(1+2m) + 6GD\mu_t + 4D\sqrt{\delta}} \\ &= \mathcal{O} \left(\sqrt{\log(m) \sum_{\tau=i}^t f_{\tau}(\mathbf{x}_{\tau}) + \log(m)} \right), \end{aligned}$$

where we denote $\mu_t = \ln(1 + (1 + \ln(1 + t))/(2e))$.

Also, Algorithm 2 ensures the bound in terms of the cumulative loss of the comparator,

$$\sum_{\tau=i}^t f_{\tau}(\mathbf{x}_{\tau}) - \sum_{\tau=i}^t f_{\tau}(\mathbf{u})$$

$$\begin{aligned}
 &\leq 4D \left(\sqrt{\ln(1+2m)} + \mu_t + 2 \right) \sqrt{L \sum_{\tau=i}^t f_{\tau}(\mathbf{u})} \\
 &\quad + (27GD + 72D^2L) \ln(1+2m) + 72D^2L\mu_t^2 + 9GD\mu_t + 6D\sqrt{\delta} + 288D^2L \quad (34) \\
 &= \mathcal{O} \left(\sqrt{\log(m) \sum_{\tau=i}^t f_{\tau}(\mathbf{u}) + \log(m)} \right).
 \end{aligned}$$

It is worth emphasizing that, the regret bound in terms of the cumulative loss of final decisions $\{\mathbf{x}_t\}$ plays a key role in proving the worst-case adaptive regret bound in Theorem 4.

The next lemma states that by the smooth and non-negative nature of loss functions, we can estimate the cumulative loss of any comparator $\mathbf{u} \in \mathcal{X}$ by the markers maintained by the problem-dependent scheduling.

Lemma 6. *Under Assumptions 1, 2, and 3, for any interval $[s_m, s_{m+1} - 1]$ determined by two consecutive intervals s_m and s_{m+1} , where we denote by s_m the m -th marker, Algorithm 2 ensures that for any comparator $\mathbf{u} \in \mathcal{X}$,*

$$\sum_{t=s_m}^{s_{m+1}-1} f_t(\mathbf{u}) \geq \frac{1}{4} C_m, \quad (35)$$

where $C_m = \mathcal{G}(m)$ is the m -th threshold with the threshold function $\mathcal{G}(\cdot)$ defined at Eq. (15).

The above lemmas rely on the unknown variable of m , which represents the number of base-learners initialized till time stamp t . The following lemma shows that m is of the same order with the cumulative loss $\sum_{\tau=1}^t f_{\tau}(\mathbf{u})$ of any comparator $\mathbf{u} \in \mathcal{X}$, owing to the construction of the problem-dependent covering intervals.

Lemma 7. *Under Assumptions 1, 2, and 3, for any interval $[i, j] \in \tilde{\mathcal{C}}$ and any $t \in [i, j]$, the variable m specified in Lemma 4 and Lemma 5 can be bounded by*

$$m \leq 1 + \frac{4}{C_1} \min_{\mathbf{u} \in \mathcal{X}} \sum_{\tau=1}^t f_{\tau}(\mathbf{u}) = \mathcal{O}(F_{[1,t]}),$$

where C_1 is a constant calculated by the threshold function as $C_1 = \mathcal{G}(1)$ defined in Eq. (15).

Moreover, for any $t \in [i, j]$ Algorithm 2 ensures,

$$\sum_{\tau=i}^t f_{\tau}(\mathbf{x}_{\tau}) - \min_{\mathbf{u} \in \mathcal{X}} \sum_{\tau=i}^t f_{\tau}(\mathbf{u}) \leq \alpha(t) + \beta(t) \sqrt{\min_{\mathbf{u} \in \mathcal{X}} \sum_{\tau=i}^t f_{\tau}(\mathbf{u})} = \mathcal{O} \left(\sqrt{F_{[i,t]} \log F_{[1,t]}} \right),$$

where

$$\begin{aligned}
 \alpha(t) &= (27GD + 72D^2L) \ln \left(3 + \frac{8}{C_1} \sum_{\tau=1}^t f_{\tau}(\mathbf{u}) \right) + 72D^2L\mu_t^2 + 9GD\mu_t + 6D\sqrt{\delta} + 288D^2L, \\
 \beta(t) &= 4D\sqrt{L} \left(\sqrt{\ln \left(3 + \frac{8}{C_1} \sum_{\tau=1}^t f_{\tau}(\mathbf{u}) \right)} + \mu_t + 2 \right),
 \end{aligned}$$

$\mu_t = \ln(1 + (1 + \ln(1 + t))/(2e))$, and $F_{[a,b]} = \min_{\mathbf{u} \in \mathcal{X}} \sum_{\tau=a}^b f_{\tau}(\mathbf{u})$ denotes the cumulative loss of the comparator within the interval $[a, b] \subseteq [T]$.

C.2 Proof of Theorem 4

Proof. The statement in Theorem 4 consists of two parts, including a small-loss bound of $\mathcal{O}(\sqrt{F_I} \log F_I \log F_T)$ and a worst-case bound of $\mathcal{O}(\sqrt{|I|} \log T)$. Below, we present the proofs of the two bounds respectively.

Small-loss regret bound. Given any interval $[r, s] \subseteq [T]$, we will identify a series of intervals I_1, \dots, I_v in the schedule $\tilde{\mathcal{C}}$ that almost covers the entire interval $[r, s]$. Then, we can use Lemma 7 to ensure low regret over these intervals. By further demonstrating that the regret on the uncovered interval can be well-controlled and that the number of intervals v is not large, we can ultimately achieve the desired bound. Below, we provide formal proof.

Recall in the algorithmic procedures, the algorithm will register a series of markers s_1, s_2, \dots . Let s_p be the smallest marker that is larger than r , and let s_q be the largest marker that is not larger than s . As a result, we have $s_{p-1} \leq r < s_p$, and $s_q \leq s < s_{q+1}$.

We bound the regret over the interval $[r, s_p - 1]$ that is not covered by the schedule as

$$\sum_{t=r}^{s_p-1} f_t(\mathbf{x}_t) - \sum_{t=r}^{s_p-1} f_t(\mathbf{u}) \leq \sum_{t=r}^{s_p-1} f_t(\mathbf{x}_t) \leq \sum_{t=s_{p-1}}^{s_p-1} f_t(\mathbf{x}_t) \leq C_{p-1} + GD. \quad (36)$$

The first and the second inequalities are by the non-negative property of loss functions. The last inequality is due to $\sum_{t=s_{p-1}}^{s_p-1} f_t(\mathbf{x}_t) \leq C_{p-1}$, which is determined by the threshold mechanism, and the fact that $f_t(\mathbf{x}_t) \leq GD$ by Assumptions 1 to 3.

Next, we focus on the interval $[s_p, s]$, which can be covered by geometric covering intervals. By Lemma 19, we can find v consecutive intervals

$$I_1 = [s_{i_1}, s_{i_2} - 1], I_2 = [s_{i_2}, s_{i_3} - 1], \dots, I_v = [s_{i_v}, s_{i_{v+1}} - 1] \in \tilde{\mathcal{C}}, \quad (37)$$

such that $i_1 = p$, $i_v \leq q < i_{v+1}$, and $v \leq \lceil \log_2(q - p + 2) \rceil$. Also notice that,

$$q < i_{v+1} \Rightarrow q + 1 \leq i_{v+1} \Rightarrow s_{q+1} - 1 \leq s_{i_{v+1}} - 1 \Rightarrow s \leq s_{i_{v+1}} - 1.$$

By Lemma 7, our algorithm has anytime regret bounds on intervals I_1 to I_v , since they belong to the covering intervals $\tilde{\mathcal{C}}$,

$$\begin{aligned} \sum_{t=s_p}^s f_t(\mathbf{x}_t) - \sum_{t=s_p}^s f_t(\mathbf{u}) &= \sum_{k=1}^{v-1} \sum_{t \in I_k} (f_t(\mathbf{x}_t) - f_t(\mathbf{u})) + \sum_{t \in [s_{i_v}, s]} (f_t(\mathbf{x}_t) - f_t(\mathbf{u})) \\ &\leq \sum_{k=1}^{v-1} \left(\alpha(s_{i_{k+1}} - 1) + \beta(s_{i_{k+1}} - 1) \sqrt{F_{I_k}} \right) + \alpha(s) + \beta(s) \sqrt{F_{[s_{i_v}, s]}} \\ &\leq \sum_{k=1}^{v-1} \left(\alpha(s) + \beta(s) \sqrt{F_{I_k}} \right) + \alpha(s) + \beta(s) \sqrt{F_{[s_{i_v}, s]}} \end{aligned} \quad (38)$$

$$\begin{aligned} &\leq v\alpha(s) + \beta(s) \sqrt{vF_{[s_p, s]}} \\ &\leq v\alpha(s) + \beta(s) \sqrt{vF_I}, \end{aligned} \quad (39)$$

where the second inequality is because $\alpha(\cdot), \beta(\cdot)$ are monotonically increasing, and the last inequality is by the non-negativity of loss functions.

Combining (36) with (39), the adaptive regret on any interval $[r, s]$ will be

$$\begin{aligned} \sum_{t=r}^s f_t(\mathbf{x}_t) - \sum_{t=r}^s f_t(\mathbf{u}) &= \sum_{t=r}^{s_p-1} (f_t(\mathbf{x}_t) - f_t(\mathbf{u})) + \sum_{t=s_p}^s (f_t(\mathbf{x}_t) - f_t(\mathbf{u})) \\ &\leq C_{p-1} + GD + v\alpha(s) + \beta(s)\sqrt{vF_I}. \end{aligned} \quad (40)$$

Furthermore, we show that C_{p-1} and v are of order $\mathcal{O}(\log F_T)$ and $\mathcal{O}(\log F_{[r,s]})$ respectively. By the definition of the time-varying threshold (see the threshold generating function Eq. (15)) and the upper bound of m in Lemma 7, the threshold can be bounded as,

$$C_{p-1} \leq (54GD + 168D^2L) \ln \left(3 + \frac{8}{C_1} F_{[1,r]} \right) + 168D^2L\mu_T^2 + 18GD\mu_T + 6D\sqrt{\delta} + 672D^2L,$$

with C_1 and μ_T defined in Lemma 7. Notice that, we treat $\mu_T = \mathcal{O}(\log \log T)$ as a constant following previous studies (Luo and Schapire, 2015; Gaillard et al., 2014; Zhao et al., 2024).

Through $[s_p, s_q - 1]$, the algorithm registers $q - p$ markers, i.e., $s_p, s_{p+1}, \dots, s_{q-1}$, then by Lemma 6 we can lower bound the cumulative loss of comparator $\mathbf{u}^* \triangleq \arg \min_{\mathbf{u} \in \mathcal{X}} \sum_{t=r}^s f_t(\mathbf{u})$ by the corresponding thresholds,

$$\sum_{t=s_p}^{s_q-1} f_t(\mathbf{u}^*) = \sum_{i=p}^{q-1} \sum_{t \in [s_i, s_{i+1}-1]} f_t(\mathbf{u}^*) \geq \frac{1}{4} \sum_{i=p}^{q-1} C_i \geq \frac{C_1}{4} (q - p),$$

where the last inequality is because the thresholds are monotonically increasing. The above inequality immediately implies that $q - p \leq \frac{4}{C_1} \sum_{t=s_p}^{s_q-1} f_t(\mathbf{u}^*) \leq \frac{4}{C_1} \sum_{t=r}^s f_t(\mathbf{u}^*)$ by the non-negativity of loss functions. We estimate the number of intervals v with Lemma 19,

$$v \leq \lceil \log_2(q - p + 2) \rceil \leq \left\lceil \log_2 \left(\frac{4}{C_1} F_{[r,s]} + 2 \right) \right\rceil = \mathcal{O}(\log F_{[r,s]}). \quad (41)$$

Combining the upper bounds of C_{p-1} and v , the adaptive regret bound in (40) as well as the definition of $\alpha(\cdot), \beta(\cdot)$ in Lemma 7 yields,

$$\begin{aligned} \sum_{t=r}^s f_t(\mathbf{x}_t) - \sum_{t=r}^s f_t(\mathbf{u}) &\leq C_{p-1} + v\alpha(s) + \beta(s)\sqrt{vF_I} + GD \\ &\leq \mathcal{O}(\log F_T) + \mathcal{O}(\log F_I \log F_T) + \mathcal{O}(\sqrt{F_I \log F_I \log F_T}) + \mathcal{O}(1) \\ &= \mathcal{O}(\sqrt{F_I \log F_I \log F_T}), \end{aligned}$$

where the last step is true as we follow the same convention in (Zhang et al., 2019) to treat the $\mathcal{O}(\log F_I \log F_T)$ as the non-leading term. We finish the proof for small-loss regret bound.

Worst-case regret bound. The above proof aims at obtaining small-loss type regret bound, and one of the key steps is to use Cauchy-Schwarz inequality to bound (38), which results in an additional $\mathcal{O}(\sqrt{\log F_{[r,s]}})$ term. Next, we show that asymptotically this extra term can be avoided thanks to the new-designed thresholds mechanism. Thus, our algorithm achieves the same worst-case adaptive regret as the best known result (Jun et al., 2017).

The primary insight for this proof lies in employing the summation for a geometric series instead of using the Cauchy-Schwarz inequality to combine the regret bounds on the intervals I_1, \dots, I_v . The crucial ingredient is to demonstrate that the worst-case regret bound on each interval scales with the number of markers within it, and that the number of these markers within each interval constitutes a geometric series.

From Lemma 5, we have that for any interval $I = [i, j]$ in problem-dependent schedule defined in (11), the adaptive regret is at most

$$\sum_{t=i}^j f_t(\mathbf{x}_t) - \sum_{t=i}^j f_t(\mathbf{u}) \leq \mathcal{O}\left(\sqrt{\log T \cdot F_{[i,j]}^{\mathbf{x}}} + \log T\right), \quad (42)$$

where we use the notation $F_{[a,b]}^{\mathbf{x}} = \sum_{t=a}^b f_t(\mathbf{x}_t)$ to denote the cumulative loss of the final decisions within the interval $[a, b] \subseteq [T]$, and we apply Lemma 7 to upper bound $m \leq \mathcal{O}(T)$ as only the worst-case behavior matters now.

For the intervals $I_k = [s_{i_k}, s_{i_{k+1}} - 1], k \in [v]$ defined in (37), we have the following facts:

$$i_{k+1} \leq 2 \cdot i_k, \quad \forall k \in [v], \quad \text{and} \quad |i_{k+1} - i_k| \leq \frac{1}{2} |i_{k+2} - i_{k+1}|, \quad \forall k \in [v-1]. \quad (43)$$

The first inequality above, which can be verified by the construction of cover defined in (11), is used to show that the time-varying thresholds do not grow too fast. The second inequality, which can be found in the proof of (Zhang et al., 2019, Lemma 11), indicates that the number of markers within each interval decreases exponentially from I_v to I_1 .

For any interval I_k with $k \in [v-1]$ in (37), our algorithm's cumulative loss within the interval can be upper bounded as

$$\begin{aligned} \sum_{t=s_{i_k}}^{s_{i_{k+1}}-1} f_t(\mathbf{x}_t) &= \sum_{a=i_k}^{i_{k+1}-1} \left(\left(\sum_{t \in [s_a, s_{a+1}-2]} f_t(\mathbf{x}_t) \right) + f_{s_{a+1}-1}(\mathbf{x}_{s_{a+1}-1}) \right) \\ &\leq \left(\sum_{a=i_k}^{i_{k+1}-1} C_a \right) + GD |i_{k+1} - i_k| \leq (GD + C_{i_{k+1}-1}) |i_{k+1} - i_k|. \end{aligned} \quad (44)$$

where the first inequality is by the threshold mechanism and the fact that $f_t(\mathbf{x}) \in [0, GD]$.

We then split a given interval $[r, s]$ into three parts to analyze, namely, the consecutive $v-1$ intervals I_1 to I_{v-1} , interval $[r, s_p - 1]$, and $[s_{i_v}, s]$, where notably the last two intervals are not fully covered by any interval in geometric covering intervals,

$$\begin{aligned} \sum_{t=r}^s f_t(\mathbf{x}_t) - \sum_{t=r}^s f_t(\mathbf{u}) &= \sum_{t=r}^{s_p-1} f_t(\mathbf{x}_t) - f_t(\mathbf{u}) + \sum_{t=s_p}^{s_{i_v}-1} f_t(\mathbf{x}_t) - f_t(\mathbf{u}) + \sum_{t=s_{i_v}}^s f_t(\mathbf{x}_t) - f_t(\mathbf{u}) \\ &= \underbrace{\sum_{t=r}^{s_p-1} f_t(\mathbf{x}_t) - f_t(\mathbf{u})}_{\text{term-(a)}} + \underbrace{\sum_{k=1}^{v-1} \sum_{t \in I_k} f_t(\mathbf{x}_t) - f_t(\mathbf{u})}_{\text{term-(b)}} + \underbrace{\sum_{t=s_{i_v}}^s f_t(\mathbf{x}_t) - f_t(\mathbf{u})}_{\text{term-(c)}}. \end{aligned} \quad (45)$$

We analyze **term-(b)** first, since it is the most intricate part in this proof. From interval I_1 to I_{v-1} , beginning with [Eq. \(42\)](#) we have

$$\begin{aligned}
 \text{term-(b)} &\leq \sum_{k=1}^{v-1} \mathcal{O} \left(\sqrt{\log T \cdot F_{I_k}^{\mathbf{x}}} + \log T \right) \\
 &\leq \sum_{k=1}^{v-1} \mathcal{O} \left(\sqrt{\log T \cdot C_{i_{v-1}} \cdot |i_{k+1} - i_k|} + \log T \right) \\
 &\leq \sum_{k=1}^{v-1} \mathcal{O} \left(\sqrt{\log T \cdot C_{i_{v-1}} \cdot \frac{|i_v - i_{v-1}|}{2^{v-1-k}}} + \log T \right) \\
 &\leq \mathcal{O} \left(v \log T + \sqrt{\log T \cdot C_{i_{v-1}}} \cdot \sum_{b=0}^{+\infty} \sqrt{\frac{|i_v - i_{v-1}|}{2^b}} \right) \\
 &\leq \mathcal{O} \left(v \log T + \sqrt{\log T \cdot C_{i_{v-1}} \cdot |i_v - i_{v-1}|} \right), \tag{46}
 \end{aligned}$$

where the second inequality is by [\(44\)](#) and together with the monotonically increasing property of thresholds, the third inequality is by the second inequality listed in [\(43\)](#), and the last inequality is by the summation of geometric sequence. We emphasize that the second inequality is upheld due to the newly-designed problem-dependent schedule mechanism. This mechanism, which monitors the cumulative loss of final decisions $\{f_t(\mathbf{x}_t)\}$, enables us to associate the $F_{I_k}^{\mathbf{x}}$ factor with the number of markers $|i_{k+1} - i_k|$ and further to apply the summation of geometric series.

In the subsequent analysis, our objective is to demonstrate that $C_{i_{v-1}} \cdot |i_v - i_{v-1}| = \mathcal{O}(|I|)$. Employing the mechanism of the time-varying threshold as defined in [Eq. \(15\)](#), we have

$$C_{i_{v-1}} = \mathcal{G}(i_v - 1) \leq \mathcal{O}(\log(i_v)).$$

Moreover, since $|i_v - i_{v-1}|$ denotes the number of markers generated by our algorithm during the interval I_{v-1} , it can be bounded above by

$$|i_v - i_{v-1}| \leq \frac{GD|I|}{C_{i_{v-1}}} = \mathcal{O} \left(\frac{GD|I|}{\log(i_{v-1})} \right).$$

This stems from the fact that the cumulative loss of the algorithm over I does not exceed $GD|I|$, and we leverage the smallest threshold $C_{i_{v-1}}$ during the interval I_{v-1} to determine the upper limit on the number of markers.

Plugging the upper bounds for $C_{i_{v-1}}$ and $|i_v - i_{v-1}|$ into [Eq. \(46\)](#), we have

$$\begin{aligned}
 \text{term-(b)} &\leq \mathcal{O} \left(v \log T + \sqrt{\log T \cdot C_{i_{v-1}} \cdot |i_v - i_{v-1}|} \right) \\
 &\leq \mathcal{O} \left(v \log T + \sqrt{\log T \cdot \log(i_v) \cdot \frac{|I|}{\log(i_{v-1})}} \right) \\
 &\leq \mathcal{O} \left(v \log T + \sqrt{\log T \cdot |I| \left(1 + \frac{1}{\log i_{v-1}} \right)} \right)
 \end{aligned}$$

$$\leq \mathcal{O}\left(\log |I| \log T + \sqrt{|I| \log T}\right).$$

The third inequality follows from the first inequality presented in (43). The concluding inequality is by the fact that $v \leq \mathcal{O}(\log F_{[r,s]}) \leq \mathcal{O}(\log |I|)$ as proved in (41). This is true because the variable v is introduced in our analysis by Lemma 19, which is independent of the worst-case analysis.

As shown in Eq. (36), we can upper bound the **term-(a)** as,

$$\mathbf{term-(a)} \leq C_{p-1} \leq \mathcal{O}(\log T).$$

Using again Lemma 7, the **term-(c)** is bounded as,

$$\mathbf{term-(c)} \leq \mathcal{O}\left(\log T + \sqrt{F_{[s_{i_v}, s]} \log T}\right) \leq \mathcal{O}\left(\log T + \sqrt{|I| \log T}\right).$$

Now we are ready to derive the worst-case adaptive regret by plugging the upper bounds from **term-(a)** to **term-(c)** into Eq. (45),

$$\begin{aligned} & \sum_{t=r}^s f_t(\mathbf{x}_t) - \sum_{t=r}^s f_t(\mathbf{u}) \\ &= \sum_{t=r}^{s_p-1} f_t(\mathbf{x}_t) - f_t(\mathbf{u}) + \sum_{k=1}^{v-1} \sum_{t \in I_k} f_t(\mathbf{x}_t) - f_t(\mathbf{u}) + \sum_{t=s_{i_v}}^s f_t(\mathbf{x}_t) - f_t(\mathbf{u}) \\ &\leq \mathcal{O}(\log T) + \mathcal{O}\left(\log |I| \log T + \sqrt{|I| \log T}\right) + \mathcal{O}\left(\log T + \sqrt{|I| \log T}\right) \\ &= \mathcal{O}\left(\sqrt{|I| \log T} + \log |I| \log T\right) = \mathcal{O}\left(\sqrt{(|I| + \log T \cdot \log^2 |I|) \log T}\right) = \mathcal{O}(\sqrt{|I| \log T}). \end{aligned}$$

The last step holds by considering the following cases.

- When the interval length is $|I| = \Theta(T^\alpha)$ with $\alpha \in (0, 1]$. Then,

$$\begin{aligned} & \mathcal{O}\left(\sqrt{(|I| + \log T \cdot \log^2 |I|) \log T}\right) \\ &= \mathcal{O}\left(\sqrt{(T^\alpha + \alpha^2 \log^3 T) \log T}\right) \\ &= \mathcal{O}\left(\sqrt{T^\alpha \log T}\right) = \mathcal{O}(\sqrt{|I| \log T}). \end{aligned}$$

- When the interval length is $|I| = \Theta(\log^\beta T)$, and note that $\beta \in [1, +\infty)$ as $|I| = \Omega(\log T)$ is the minimum order to ensure the adaptive regret to be non-trivial. Then,

$$\begin{aligned} & \mathcal{O}\left(\sqrt{(|I| + \log T \cdot \log^2 |I|) \log T}\right) \\ &= \mathcal{O}\left(\sqrt{(\log^\beta T + \beta^2 \log T \cdot (\log \log T)^2) \log T}\right) \\ &= \mathcal{O}\left(\sqrt{(\log^\beta T + \beta^2 \log T) \log T}\right) \\ &= \mathcal{O}\left(\sqrt{\log^\beta T \log T}\right) = \mathcal{O}(\sqrt{|I| \log T}). \end{aligned}$$

Hence we finish the proof for the worst-case adaptive regret bound. Combining both small-loss bound and the worst-case safety guarantee, we complete the proof of Theorem 4. \square

C.3 Proof of Lemma 4

Proof. First we introduce some useful variables to help us prove the adaptivity of Adapt-ML-Prod under sleeping-expert setting. Similar to the proof technique of Daniely et al. (2015), for any interval $[i, j] \in \mathcal{C}$ in the geometric covers defined in (11), on which we suppose the m -th base-learner is active, we define the following pseudo-weight for the m -th base-learner,

$$\tilde{w}_{\tau,m} = \begin{cases} 0 & \tau < i, \\ 1 & \tau = i, \\ (\tilde{w}_{\tau-1,m}(1 + \eta_{\tau-1,m}(\hat{\ell}_{\tau-1} - \ell_{\tau-1,m})))^{\frac{\eta_{\tau,m}}{\eta_{\tau-1,m}}} & i < \tau \leq j+1, \\ \tilde{w}_{j+1,m} & \tau > j+1. \end{cases}$$

In addition, we use $\tilde{W}_t = \sum_{k \in [T]} \tilde{w}_{t,k}$ to denote the summation of pseudo-weights for all possible base-learners up to time t . As for the problem-dependent geometric covers, in the worst case there are at most T base-learners generated, we use $[T]$ to denote the indexes for all the base-learners. Notice that the pseudo-weight \tilde{w}_t is defined as 0 for asleep base-learners till time t , so we can include all possible ones safely in the definition even though they are not generated in practical implementations of the algorithm.

Below, we use the potential argument (Gaillard et al., 2014) to prove the desired result. Specifically, we establish the regret bound by lower and upper bounding the quantity $\ln \tilde{W}_{t+1}$.

Lower bound of $\ln \tilde{W}_{t+1}$. We claim that for $t \in [i, j]$ it holds that

$$\ln \tilde{w}_{t+1,m} \geq \eta_{t+1,m} \sum_{\tau=i}^t (r_{\tau,m} - \eta_{\tau,m} r_{\tau,m}^2).$$

We prove the above inequality by induction on t . When $t = i$, by definition,

$$\ln \tilde{w}_{i+1,m} = \frac{\eta_{i+1,m}}{\eta_{i,m}} \ln(1 + \eta_m r_{i,m}) \geq \frac{\eta_{i+1,m}}{\eta_{i,m}} (\eta_m r_{i,m} - \eta_m^2 r_{i,m}^2) = \eta_{i+1,m} (r_{i,m} - \eta_m r_{i,m}^2),$$

where the inequality is because of $\ln(1+x) \geq x - x^2, \forall x \geq -1/2$.

Suppose the statement holds for $\ln \tilde{w}_{t,m}$, then we proceed to check the situation for $t+1$ round as follows. Indeed,

$$\begin{aligned} \ln \tilde{w}_{t+1,m} &= \frac{\eta_{t+1,m}}{\eta_{t,m}} (\ln \tilde{w}_{t,m} + \ln(1 + \eta_{t,m} r_{t,m})) \\ &\geq \frac{\eta_{t+1,m}}{\eta_{t,m}} (\ln \tilde{w}_{t,m} + \eta_{t,m} r_{t,m} - \eta_{t,m}^2 r_{t,m}^2) \\ &= \frac{\eta_{t+1,m}}{\eta_{t,m}} \ln \tilde{w}_{t,m} + \eta_{t+1,m} (r_{t,m} - \eta_{t,m} r_{t,m}^2) \\ &\geq \frac{\eta_{t+1,m}}{\eta_{t,m}} \left(\eta_{t,m} \sum_{\tau=i}^{t-1} (r_{\tau,m} - \eta_{\tau,m} r_{\tau,m}^2) \right) + \eta_{t+1,m} (r_{t,m} - \eta_{t,m} r_{t,m}^2) \end{aligned}$$

$$= \eta_{t+1,m} \sum_{\tau=i}^t (r_{\tau,m} - \eta_{\tau,m} r_{\tau,m}^2). \quad (47)$$

Then, as $\tilde{w}_{t+1,m}$ is positive for any m -th base-learner, we have $\ln \widetilde{W}_{t+1} \geq \ln \tilde{w}_{t+1,m}$. Thus by (47) we obtain the desired lower bound of $\ln \widetilde{W}_{t+1}$.

Upper bound of $\ln \widetilde{W}_{t+1}$. By the construction of the geometric covers as specified in Eq. (11), we know that there will be at most $2m$ base-learners initialized for the m -th base-learner active on interval $[i, j]$ till her end. This is because m -th base-learner is initialized when m -th marker is recorded, and she will expire before the moment when $2m$ -th marker is recorded, as demonstrated by the construction of cover defined in Eq. (11). Owing to this property, we have $\widetilde{W}_{t+1} = \sum_{k \in [2m]} \tilde{w}_{t+1,k}$ as others' pseudo-weight equals to 0 by definition. So we can upper bound \widetilde{W}_{t+1} as,

$$\begin{aligned} \widetilde{W}_{t+1} &= \sum_{k \in [2m]} \tilde{w}_{t+1,k} = \sum_{k \in [2m]: i_k = t+1} \tilde{w}_{t+1,k} + \sum_{k \in [2m]: i_k \leq t} \tilde{w}_{t+1,k} \\ &= \mathbf{1}\{\text{new alg. at } t+1\} + \sum_{k \in [2m]: i_k \leq t} \tilde{w}_{t+1,k}, \end{aligned} \quad (48)$$

where we denote by $[i_k, j_k] \in \tilde{\mathcal{C}}$ the active time for k -th base-learner.

For the second term in (48), we have

$$\begin{aligned} \sum_{k: i_k \leq t} \tilde{w}_{t+1,k} &= \sum_{k \in [2m]: t \in [i_k, j_k]} \tilde{w}_{t+1,k} + \sum_{k \in [2m]: t > j_k} \tilde{w}_{t+1,k} \\ &= \sum_{k \in [2m]: t \in [i_k, j_k]} \tilde{w}_{t+1,k} + \sum_{k \in [2m]: t > j_k} \tilde{w}_{t,k} \\ &\leq \sum_{k \in [2m]: t \in [i_k, j_k]} \tilde{w}_{t,k} (1 + \eta_{t,k} r_{t,k}) + \frac{1}{e} \left(\frac{\eta_{t,k}}{\eta_{t+1,k}} - 1 \right) + \sum_{k \in [2m]: t > j_k} \tilde{w}_{t,k} \\ &= \underbrace{\widetilde{W}_t + \sum_{k \in [2m]: t \in [i_k, j_k]} \eta_{t,k} \tilde{w}_{t,k} r_{t,k}}_{=0} + \sum_{k \in [2m]: t \in [i_k, j_k]} \frac{1}{e} \left(\frac{\eta_{t,k}}{\eta_{t+1,k}} - 1 \right), \end{aligned} \quad (49)$$

where the first equality holds by the definition of $\tilde{w}_{t+1,k}$, the second inequality is by the updating rule of $\tilde{w}_{t+1,k}$ and Lemma (26), and the second term in the last equality equals to 0 due to the weight update rule in (14) and the fact of $\tilde{w}_{t,k} = w_{t,k}$ for any $t \in [i_k, j_k]$.

Combining (48), (49) and by induction, we obtain the following upper bound:

$$\widetilde{W}_{t+1} \leq 1 + 2m + \frac{1}{e} \sum_{k \in [2m]} \sum_{\tau=i_k}^{\min\{t, j_k\}} \left(\frac{\eta_{\tau,k}}{\eta_{\tau+1,k}} - 1 \right). \quad (50)$$

We now turn to analyze the third term in (49). Indeed, Gaillard et al. (2014) have analyzed it under the static regret measure. For the sake of completeness, we present the proof with our notations. For any $k \in [2m]$, for any $\tau \in [i_k, \min\{t, j_k\}]$, the relationship between $\eta_{\tau,k}$ and $\eta_{\tau+1,k}$ can be considered as three cases,

- $\eta_{\tau,k} = \eta_{\tau+1,k} = 1/2$,
- $\eta_{\tau+1,k} = \sqrt{\gamma_k/(1 + \sum_{u=i_k}^{\tau} r_{u,k}^2)} < \eta_{\tau,k} = \frac{1}{2}$,
- $\eta_{\tau+1,k} \leq \eta_{\tau,k} < 1/2$.

In all cases, the ratio $\eta_{\tau,k}/\eta_{\tau+1,k} - 1$ is at most

$$\begin{aligned}
 \sum_{\tau=i_k}^{\min\{t,j_k\}} \left(\frac{\eta_{\tau,k}}{\eta_{\tau+1,k}} - 1 \right) &\leq \sum_{\tau=i_k}^{\min\{t,j_k\}} \left(\sqrt{\frac{1 + \sum_{u=i_k}^{\tau} r_{u,k}^2}{1 + \sum_{u=i_k}^{\tau-1} r_{u,k}^2}} - 1 \right) \\
 &= \sum_{\tau=i_k}^{\min\{t,j_k\}} \left(\sqrt{\frac{r_{\tau,k}^2}{1 + \sum_{u=i_k}^{\tau-1} r_{u,k}^2}} + 1 - 1 \right) \\
 &\leq \frac{1}{2} \sum_{\tau=i_k}^{\min\{t,j_k\}} \frac{r_{\tau,k}^2}{1 + \sum_{u=i_k}^{\tau-1} r_{u,k}^2} \\
 &\leq \frac{1}{2} \left(1 + \ln \left(1 + \sum_{u=i_k}^{\min\{t,j_k\}} r_{u,k}^2 \right) \right) - \ln(1) \\
 &\leq \frac{1}{2} (1 + \ln(1+t)), \tag{51}
 \end{aligned}$$

where the second inequality uses $\sqrt{1+x} \leq 1 + x/2$ and the third inequality follows from Lemma 22 with the choice of $f(x) = 1/x$. Substituting (51) into (50), we get

$$\widetilde{W}_{t+1} \leq 1 + 2m + \frac{m}{e} (1 + \ln(1+t)) \leq (1 + 2m) \left(1 + \frac{1}{2e} (1 + \ln(1+t)) \right). \tag{52}$$

Further taking the logarithm over the above inequality gives the upper bound of $\ln \widetilde{W}_{t+1}$.

Upper bound of meta-regret. Now, we can lower bound and upper bound $\ln \widetilde{W}_{t+1}$ by (47) and (52). Then, rearranging the terms yields the upper bound of scaled meta-regret,

$$\begin{aligned}
 \sum_{\tau=i}^t r_{\tau,m} &\leq \sum_{\tau=i}^t \eta_{\tau,m} r_{\tau,m}^2 + \frac{\ln(1+2m) + \mu_t}{\eta_{t+1,m}} \\
 &\leq 2\sqrt{\gamma_i} \sqrt{1 + \sum_{\tau=i}^t r_{\tau,i}^2} + \frac{\ln(1+2m) + \mu_t}{\eta_{t+1,m}} \\
 &\leq \frac{\ln(1+2m) + \mu_t + 2\gamma_m}{\sqrt{\gamma_m}} \sqrt{1 + \sum_{\tau=i}^t r_{\tau,m}^2} + 2\ln(1+2m) + 4\gamma_m + 2\mu_t \\
 &\leq \left(3\sqrt{\ln(1+2m)} + \mu_t \right) \sqrt{1 + \sum_{\tau=i}^t r_{\tau,m}^2} + 6\ln(1+2m) + 2\mu_t, \tag{53}
 \end{aligned}$$

where we denote $\mu_t = \ln(1 + (1 + \ln(1 + t))/(2e))$. The second inequality is by Lemma 22 and choose $f(x) = 1/\sqrt{x}$. The last inequality is by the choice of $\sqrt{\gamma_m} = \sqrt{\ln(1 + 2m)} \geq \sqrt{\ln(3)} \geq 1$. As for the third inequality, there are two cases to be considered:

- when $\sqrt{1 + \sum_{\tau=i}^t r_{\tau,m}^2} > 2\sqrt{\gamma_m}$, we have that (53) is at most $2\sqrt{\gamma_m} \sqrt{1 + \sum_{\tau=i}^t r_{\tau,m}^2} + \frac{\ln(1+2m) + \mu_t}{\sqrt{\gamma_m}} \sqrt{1 + \sum_{\tau=i}^t r_{\tau,m}^2}$.
- when $\sqrt{1 + \sum_{\tau=i}^t r_{\tau,m}^2} \leq 2\sqrt{\gamma_m}$, we have that $\eta_{t+1,m} = 1/2$ and (53) is at most $2\ln(1 + 2m) + 4\gamma_m + 2\mu_t$.

Taking both cases into account implies the desired inequality.

Finally, we end the proof by evaluating the meta-regret in terms of the surrogate loss.

$$\begin{aligned}
& \sum_{\tau=i}^t \langle \nabla g_t(\mathbf{y}_\tau), \mathbf{y}_\tau - \mathbf{y}_{\tau,m} \rangle \\
&= 2GD \cdot \sum_{\tau=i}^t r_{\tau,m} \\
&\leq 2GD \left(3\sqrt{\ln(1 + 2m)} + \mu_t \right) \sqrt{1 + \sum_{\tau=i}^t r_{\tau,m}^2 + 12GD \ln(1 + 2m) + 4GD\mu_t} \\
&\leq \left(3\sqrt{\ln(1 + 2m)} + \mu_t \right) \sqrt{\sum_{\tau=i}^t \langle \nabla g_\tau(\mathbf{y}_\tau), \mathbf{y}_\tau - \mathbf{y}_{\tau,m} \rangle^2 + 18GD \ln(1 + 2m) + 6GD\mu_t} \\
&\leq \left(3\sqrt{\ln(1 + 2m)} + \mu_t \right) \sqrt{\sum_{\tau=i}^t 4D^2 \|\nabla g_\tau(\mathbf{y}_\tau)\|_2^2 + 18GD \ln(1 + 2m) + 6GD\mu_t} \\
&\leq 2D \left(3\sqrt{\ln(1 + 2m)} + \mu_t \right) \sqrt{\sum_{\tau=i}^t \|\nabla f_\tau(\mathbf{x}_\tau)\|_2^2 + 18GD \ln(1 + 2m) + 6GD\mu_t} \\
&\leq 4D \left(3\sqrt{\ln(1 + 2m)} + \mu_t \right) \sqrt{L \sum_{\tau=i}^t f_t(\mathbf{x}_t) + 18GD \ln(1 + 2m) + 6GD\mu_t},
\end{aligned}$$

where the second inequality is true because $1 \leq \sqrt{\ln(1 + 2m)} \leq \ln(1 + 2m)$ holds for any $m \geq 1$, the third inequality is by Cauchy-Schwarz inequality, the forth inequality is by Theorem 1 and the last inequality is due to the self-bounding property of smooth and non-negative functions (see Lemma 20). \square

C.4 Proof of Lemma 5

Proof. We start the proof by decomposing the adaptive regret into meta-regret and base-regret in terms of the surrogate loss by Theorem 1,

$$\sum_{\tau=i}^t f_\tau(\mathbf{x}_\tau) - \sum_{\tau=i}^t f_\tau(\mathbf{u}) \leq \sum_{\tau=i}^t g_\tau(\mathbf{x}_\tau) - \sum_{\tau=i}^t g_\tau(\mathbf{u}) \leq \sum_{\tau=i}^t \langle \nabla g_\tau(\mathbf{y}_\tau), \mathbf{y}_\tau - \mathbf{u} \rangle$$

$$= \underbrace{\sum_{\tau=i}^t \langle \nabla g_\tau(\mathbf{y}_\tau), \mathbf{y}_\tau - \mathbf{y}_{\tau,m} \rangle}_{\text{meta-regret}} + \underbrace{\sum_{\tau=i}^t \langle \nabla g_\tau(\mathbf{y}_\tau), \mathbf{y}_{\tau,m} - \mathbf{u} \rangle}_{\text{base-regret}}, \quad (54)$$

where our analysis will be performed by tracking the m -th base-learner, whose corresponding active interval is exactly the analyzed one. Our analysis is satisfied to any interval since there is always a base-learner active on it ought to our algorithm design.

Upper bound of base-regret. Since the base-algorithm (SOGD) guarantees anytime regret, direct application of Lemma 16 with the assumption of surrogate domain \mathcal{Y} can upper bound the base-regret,

$$\sum_{\tau=i}^t \langle \nabla g_\tau(\mathbf{y}_\tau), \mathbf{y}_{\tau,m} - \mathbf{u} \rangle \leq 4D \sqrt{\delta + \sum_{\tau=i}^t \|\nabla g_\tau(\mathbf{y}_\tau)\|_2^2} \leq 8D \sqrt{L \sum_{\tau=i}^t f_\tau(\mathbf{x}_\tau) + 4D\sqrt{\delta}}, \quad (55)$$

where we skip some steps for transforming $\|\nabla g_\tau(\mathbf{y}_\tau)\|_2^2$ into $4Lf_\tau(\mathbf{x}_\tau)$. The similar arguments can be found in the proof of Theorem 3.

Upper bound of meta-regret. By Lemma 4, we can upper bound the meta-regret as

$$\begin{aligned} & \sum_{\tau=i}^t \langle \nabla g_\tau(\mathbf{y}_\tau), \mathbf{y}_\tau - \mathbf{y}_{\tau,m} \rangle \\ & \leq 4D \left(3\sqrt{\ln(1+2m)} + \mu_t \right) \sqrt{L \sum_{\tau=i}^t f_\tau(\mathbf{x}_\tau) + 18GD \ln(1+2m) + 6GD\mu_t}, \end{aligned} \quad (56)$$

where we denote $\mu_t = \ln(1 + (1 + \ln(1+t))/(2e))$.

Upper bound of adaptive regret. Substituting (55), (56) into (54) obtains

$$\begin{aligned} & \sum_{\tau=i}^t f_\tau(\mathbf{x}_\tau) - \sum_{\tau=i}^t f_\tau(\mathbf{u}) \\ & \leq 4D \left(3\sqrt{\ln(1+2m)} + \mu_t + 2 \right) \sqrt{L \sum_{\tau=i}^t f_\tau(\mathbf{x}_\tau) + 18GD \ln(1+2m) + 6GD\mu_t + 4D\sqrt{\delta}} \\ & = \mathcal{O} \left(\sqrt{\log(m) \cdot \sum_{\tau=i}^t f_\tau(\mathbf{x}_\tau) + \log(m)} \right), \end{aligned} \quad (57)$$

which proves the first part of the results.

Furthermore, by applying the standard technical lemma presented in Lemma 25, we can convert the cumulative loss of final decisions in the above regret bound, $\sum_{\tau=i}^t f_\tau(\mathbf{x}_\tau)$, into the comparator's cumulative loss, $\sum_{\tau=i}^t f_\tau(\mathbf{u})$,

$$\sum_{\tau=i}^t f_\tau(\mathbf{x}_\tau) - \sum_{\tau=i}^t f_\tau(\mathbf{u})$$

$$\begin{aligned}
 &\leq 4D\sqrt{L}\left(\sqrt{\ln(1+2m)} + \mu_t + 2\right)\sqrt{\sum_{\tau=i}^t f_{\tau}(\mathbf{u}) + 18GD\ln(1+2m) + 6GD\mu_t + 4D\sqrt{\delta}} \\
 &\quad + 18GD\ln(1+2m) + 6GD\mu_t + 4D\sqrt{\delta} + 16D^2L\left(\sqrt{\ln(1+2m)} + \mu_t + 2\right)^2 \\
 &\leq 4D\sqrt{L}\left(\sqrt{\ln(1+2m)} + \mu_t + 2\right)\sqrt{\sum_{\tau=i}^t f_{\tau}(\mathbf{u})} \\
 &\quad + 27GD\ln(1+2m) + 9GD\mu_t + 6D\sqrt{\delta} + 24D^2L\left(\sqrt{\ln(1+2m)} + \mu_t + 2\right)^2 \\
 &\leq 4D\sqrt{L}\left(\sqrt{\ln(1+2m)} + \mu_t + 2\right)\sqrt{\sum_{\tau=i}^t f_{\tau}(\mathbf{u})} \\
 &\quad + (27GD + 72D^2L)\ln(1+2m) + 72D^2L\mu_t^2 + 9GD\mu_t + 6D\sqrt{\delta} + 288D^2L.
 \end{aligned}$$

The second inequality makes use of $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ and $\sqrt{ab} \leq (a+b)/2$. The last inequality holds by $(a+b+c)^2 \leq 3(a^2+b^2+c^2)$. \square

C.5 Proof of Lemma 6

Proof. For interval $[s_m, s_{m+1}-1]$, there must exist an interval $[i, j] \in \tilde{\mathcal{C}}$ such that $[s_m, s_{m+1}-1] \subseteq [i, j]$ with $i = s_m$. Therefore, we can apply Eq. (34) presented in Lemma 5 to upper bound the regret during $[s_m, s_{m+1}-1]$ for any comparator $\mathbf{u} \in \mathcal{X}$,

$$\begin{aligned}
 \sum_{t=s_m}^{s_{m+1}-1} f_t(\mathbf{x}_t) - \sum_{t=s_m}^{s_{m+1}-1} f_t(\mathbf{u}) &\leq 4D\left(\sqrt{\ln(1+2m)} + \mu_T + 2\right)\sqrt{L\sum_{t=i}^j f_t(\mathbf{u})} \\
 &\quad + (27GD + 72D^2L)\ln(1+2m) + 72D^2L\mu_T^2 + 9GD\mu_T \\
 &\quad + 6D\sqrt{\delta} + 288D^2L,
 \end{aligned} \tag{58}$$

where we use the monotonically increasing property that $\mu_{s_{m+1}-1} \leq \mu_T$.

Incorporating some basic inequalities, specifically $ab \leq a^2/4 + b^2$ and $(a+b+c)^2 \leq 3(a^2+b^2+c^2)$, we can isolate $\sum_{t=s_m}^{s_{m+1}-1} f_t(\mathbf{u})$ from the square root in (58):

$$\begin{aligned}
 \sum_{t=s_m}^{s_{m+1}-1} f_t(\mathbf{x}_t) &\leq 2 \sum_{t=s_m}^{s_{m+1}-1} f_t(\mathbf{u}) + (27GD + 84D^2L)\ln(1+2m) \\
 &\quad + 84D^2L\mu_T^2 + 9GD\mu_T + 3D\sqrt{\delta} + 336D^2L \\
 &= 2 \sum_{t=s_m}^{s_{m+1}-1} f_t(\mathbf{u}) + \frac{1}{2}\mathcal{G}(m),
 \end{aligned} \tag{59}$$

where the equality is by the definition in Eq. (15).

By the problem-dependent schedule mechanism, as stated in Lines 7, the cumulative loss between $[i, j]$ exceeds the threshold $C_m = \mathcal{G}(m)$, i.e., $\sum_{t=s_m}^{s_{m+1}-1} f_t(\mathbf{x}_t) \geq C_m$. Therefore,

together with Eq. (59), we can lower bound the cumulative loss of comparator \mathbf{u} as,

$$\sum_{t=s_m}^{s_{m+1}-1} f_t(\mathbf{u}) \geq \frac{1}{2} \left(\sum_{t=s_m}^{s_{m+1}-1} f_t(\mathbf{x}_t) - \frac{1}{2} \mathcal{G}(m) \right) \geq \frac{1}{2} \left(C_m - \frac{1}{2} \mathcal{G}(m) \right) = \frac{1}{4} C_m,$$

where the last equality is by the definition of threshold. \square

C.6 Proof of Lemma 7

Proof. We assume that the m -th base-learner is initialized at the beginning of interval $[i, j] \in \tilde{\mathcal{C}}$, in other words, $i = s_m$, where s_m denotes the m -th marker. Before time stamp i , the schedule has registered $m - 1$ markers, i.e., from s_1 to s_{m-1} . By Lemma 6, we can lower bound the cumulative loss of any comparator $\mathbf{u} \in \mathcal{X}$ as,

$$\sum_{\tau=s_1}^{s_m-1} f_\tau(\mathbf{u}) \geq \frac{1}{4} \sum_{k=1}^{m-1} C_k \geq \frac{C_1}{4} (m-1).$$

The second inequality holds since C_k is monotonically increasing with respect to its index, see the threshold generating function in Eq. (15).

Therefore, rearranging the above inequalities gives upper bound of quantity m ,

$$m \leq 1 + \frac{4}{C_1} \sum_{\tau=s_1}^{s_m-1} f_\tau(\mathbf{u}) \leq 1 + \frac{4}{C_1} \sum_{\tau=1}^t f_\tau(\mathbf{u}),$$

where the last inequality makes use of the non-negative assumption on loss functions. Choosing $\mathbf{u} \in \arg \min_{\mathbf{a} \in \mathcal{X}} \sum_{\tau=1}^t f_\tau(\mathbf{a})$ as the minimizer for the given interval $[i, j]$ gives

$$m \leq 1 + \frac{4}{C_1} \min_{\mathbf{u} \in \mathcal{X}} \sum_{\tau=1}^t f_\tau(\mathbf{u}). \quad (60)$$

Plugging Eq. (60) into Eq. (34) in Lemma 5, we prove the second result in this lemma,

$$\sum_{\tau=i}^t f_\tau(\mathbf{x}_\tau) - \min_{\mathbf{u} \in \mathcal{X}} \sum_{\tau=i}^t f_\tau(\mathbf{u}) \leq \alpha(t) + \beta(t) \sqrt{\min_{\mathbf{u} \in \mathcal{X}} \sum_{\tau=i}^t f_\tau(\mathbf{u})} = \mathcal{O} \left(\sqrt{F_{[i,t]} \log F_{[1,t]}} \right),$$

where

$$\begin{aligned} \alpha(t) &= (27GD + 72D^2L) \ln \left(3 + \frac{8}{C_1} \sum_{\tau=1}^t f_\tau(\mathbf{u}) \right) + 72D^2L\mu_t^2 + 9GD\mu_t + 6D\sqrt{\delta} + 288D^2L, \\ \beta(t) &= 4D\sqrt{L} \left(\sqrt{\ln \left(3 + \frac{8}{C_1} \sum_{\tau=1}^t f_\tau(\mathbf{u}) \right)} + \mu_t + 2 \right), \end{aligned}$$

with $\mu_t = \ln(1 + (1 + \ln(1 + t))/(2e))$. \square

Appendix D. Omitted Details for Interval Dynamic Regret Minimization

In this section, we first discuss the method of Cutkosky (2020) in Appendix D.1, another algorithm designed for optimizing interval dynamic regret in the literature, and present a projection-efficient variant. Then, we provide the proof of Theorem 5, the small-loss interval dynamic regret bound with 1 projection complexity. Key lemmas supporting this proof are detailed in Appendix D.2, with further key proofs found in Appendix D.3.

D.1 Projection-Efficient Variant for the Method of Cutkosky (2020)

As discussed in the main text, Cutkosky (2020) proposes an algorithm to optimize interval dynamic regret, which significantly differs in its algorithmic framework from the meta-base structure of the AOA algorithm (Zhang et al., 2020). In this appendix, we further demonstrate how our reduction mechanism can be applied to achieve a projection-efficient variant of this algorithm. Additionally, we provide a computational comparison with our algorithm designed for optimizing interval dynamic regret (i.e., Algorithm 3).

Reviewing the method of Cutkosky (2020). Unlike the meta-base type algorithm (such as AOA (Zhang et al., 2020)) which combines decisions using a weighted vector $\mathbf{p}_t \in \Delta_N$ from the meta-algorithm as $\mathbf{x}_t = \sum_{i=1}^N p_{t,i} \mathbf{x}_{t,i}$, Cutkosky (2020) aggregates the decisions from base-learners *sequentially*, submitting the decision as $\mathbf{x}_t = \sum_{i=1}^N \mathbf{x}_{t,i}$ with the number of base-learners $N = \mathcal{O}(\log T)$. Their method relies on the specific beneficial properties of parameter-free online learners (Cutkosky and Orabona, 2018). Specifically, with $\mathbf{x}_t^{(k-1)} = \sum_{i=1}^{k-1} \mathbf{x}_{t,i}$ representing the aggregated decisions of the first $k-1$ base-learners, the k -th base-learner then makes a decision $(\mathbf{y}_{t,k}, z_{t,k})$ on a $(d+1)$ -dimensional lifted domain $V_t^{(k)}$ defined as

$$V_t^{(k)} = \{(\mathbf{y}, z) \in \mathbb{R}^d \times [0, 1] \mid \mathbf{x}_t^{(k-1)} + \mathbf{y} - z \cdot \mathbf{x}_t^{(k-1)} \in \mathcal{X}\} \subseteq \mathbb{R}^{d+1}. \quad (61)$$

The decision of the k -th base-learner is $\mathbf{x}_{t,k} = \mathbf{y}_{t,k} - z_{t,k} \cdot \mathbf{x}_t^{(k-1)}$. The projection onto $V_t^{(k)}$ ensures that the updated decision $\mathbf{x}_t^{(k)} = \mathbf{x}_t^{(k-1)} + \mathbf{x}_{t,k}$ remains within the feasible domain \mathcal{X} . Generally, projecting onto $V_t^{(k)}$ is not simpler than projecting onto the original feasible domain \mathcal{X} . Moreover, each round requires $\mathcal{O}(\log T)$ projections onto this lifted domain, significantly increasing computational costs.

Projection-Efficient Variant. By applying our reduction method to the approach of Cutkosky (2020), we can reduce the computation of the projection. Specifically, as discussed in Section 2.2, we can construct a surrogate domain $\mathcal{Y} \supseteq \mathcal{X}$ with a simpler structure, e.g., \mathcal{Y} is often chosen as an Euclidean ball. Then, for $k \in [N-1]$, we allow the base-learners to project onto the lifted domain constructed over \mathcal{Y} :

$$\widehat{V}_t^{(k)} = \{(\mathbf{y}, z) \in \mathbb{R}^d \times [0, 1] \mid \mathbf{x}_t^{(k-1)} + \mathbf{y} - z \cdot \mathbf{x}_t^{(k-1)} \in \mathcal{Y}\}. \quad (62)$$

If we choose \mathcal{Y} to be an Euclidean ball, the corresponding lifted domain $\widehat{V}_t^{(k)}$ becomes a cylinder, which simplifies the projection for the base-learners. In this approach, we only require the N -th base-learner to project her decision onto the lifted domain constructed over \mathcal{X} , to ensure the final submitted decision $\mathbf{x}_t = \mathbf{x}_t^{(N)} = \sum_{i=1}^N \mathbf{x}_{t,i} \in \mathcal{X}$ to lie within \mathcal{X} .

After submitting \mathbf{x}_t , we send the gradient $\mathbf{g}_t = \nabla f_t(\mathbf{x}_t)$ to the method of Cutkosky (2020) with the aforementioned revision. Thus, our reduction mechanism can improve efficiency by requiring only 1 projection onto the complex lifted domain associated with \mathcal{X} (as defined in (61) for $k = N$), while performing $\mathcal{O}(\log T)$ projections onto simpler lifted domains (as defined in (62) for $k = 1, \dots, N - 1$).

Comparisons. Even with our reduction mechanism applied, the variant of Cutkosky (2020)’s method still necessitates at least 1 projection onto the complex lifted domain associated with \mathcal{X} , which is more complicated than the simple projection onto \mathcal{X} required by our method. Furthermore, although performing $\mathcal{O}(\log T)$ projections onto simpler lifted domains is easy and feasible, it is still more complex than the simple rescaling operation our method uses for projections onto the (surrogate) Euclidean ball.

D.2 Key Lemmas

This part collects several key lemmas for proving small-loss interval dynamic regret (Theorem 5). The first lemma shows that our base-algorithm enjoys an anytime dynamic regret.

Lemma 8. *Under Assumptions 1, 2, and 3, setting the step size pool as $\mathcal{H} = \{\eta_j = 2^{j-1} \sqrt{5D^2/(1+8LGDT)} \mid j \in [N]\}$ with $N = \lceil 2^{-1} \log_2((5D^2+2D^2T)(1+8LGDT)/(5D^2)) \rceil + 1$ and $\varepsilon_{i,\tau} = \sqrt{(\ln N)/(1+D^2 \sum_{s=i}^{\tau-1} \|\nabla g_s(\mathbf{y}_s)\|_2^2)}$, assume that the m -th base-learner described in Algorithm 3 is initialized at the beginning of an interval $[i, j] \in \tilde{\mathcal{C}}$, then this base-learner ensures the following dynamic regret*

$$\begin{aligned} & \sum_{\tau=i}^t \langle \nabla g_\tau(\mathbf{y}_\tau), \mathbf{y}_{\tau,m} - \mathbf{u}_\tau \rangle \\ & \leq 6D \sqrt{L(\ln N) F_{[i,t]}^\mathbf{x}} + 3 \sqrt{2(5D^2 + 2DP_{[i,t]}) (1 + LF_{[i,t]}^\mathbf{x})} + \frac{(6 + G^2 D^2) \sqrt{\ln N}}{2} \\ & = \mathcal{O} \left(\sqrt{F_{[i,t]}^\mathbf{x} (1 + P_{[i,t]})} \right), \end{aligned}$$

which holds for any time stamp $t \in [i, j]$ and any comparators sequence $\mathbf{u}_i, \dots, \mathbf{u}_t \in \mathcal{X}$. In above, $\mathbf{y}_{\tau,m}$ denotes the decision of the m -th base-learner at time stamp τ , $F_{[i,t]}^\mathbf{x} = \sum_{\tau=i}^t f_\tau(\mathbf{x}_\tau)$ denotes the cumulative loss of the decisions and $P_{[i,t]} = \sum_{\tau=i+1}^t \|\mathbf{u}_\tau - \mathbf{u}_{\tau-1}\|_2$ is the path-length of comparators within the interval $[i, t]$.

Combining the above lemma with the analysis of the meta-algorithm Adapt-ML-Prod (see Lemma 4), we can obtain the interval dynamic regret for intervals in the problem-dependent geometric covering intervals $\tilde{\mathcal{C}}$.

Lemma 9. *Under Assumptions 1, 2, and 3, for any interval $I = [i, j] \in \tilde{\mathcal{C}}$ in the geometric covers defined in Eq. (11), at the beginning of which we assume the m -th base-learner is initialized, Algorithm 3 ensures*

$$\sum_{\tau=i}^t f_\tau(\mathbf{x}_\tau) - f_\tau(\mathbf{u}_\tau)$$

$$\begin{aligned}
 &\leq \left(12D\sqrt{\ln(1+2m)} + 4D\mu_t + 6D\sqrt{\ln N} + 3\sqrt{2(5D^2 + 2DP_{[i,t]})}\right) \sqrt{LF_{[i,t]}^{\mathbf{u}}} \\
 &\quad + \frac{3L}{2} \left(12D\sqrt{\ln(1+2m)} + 4D\mu_t + 6D\sqrt{\ln N} + 3\sqrt{2(5D^2 + 2DP_{[i,t]})}\right)^2 \\
 &\quad + 27GD\ln(1+2m) + 9GD\mu_t + \frac{9}{2}\sqrt{2(5D^2 + 2DP_{[i,t]})} + \frac{3(6 + G^2D^2)\sqrt{\ln N}}{4} \quad (63) \\
 &= \mathcal{O}\left(\sqrt{F_{[i,t]}^{\mathbf{u}}(P_{[i,t]} + \log m)} + P_{[i,t]}\right),
 \end{aligned}$$

which holds for any time stamp $t \in [i, j]$ and any comparators sequence $\mathbf{u}_i, \dots, \mathbf{u}_t \in \mathcal{X}$. In above, $F_{[i,t]}^{\mathbf{u}} = \sum_{\tau=i}^t f_{\tau}(\mathbf{u}_{\tau})$ denotes the cumulative loss of the comparators and $P_{[i,t]} = \sum_{\tau=i+1}^t \|\mathbf{u}_{\tau} - \mathbf{u}_{\tau-1}\|_2$ is the path-length of comparators within the interval $[i, t]$.

As an analog to Lemma 6, with the components at hand, we would like to estimate the cumulative loss $\sum_{t \in I} f_t(\mathbf{u}_t)$ in terms of the thresholds. However, our analysis shows that, we cannot estimate the cumulative loss alone, instead with the path length together, $F_I^{\mathbf{u}} + P_I = \sum_{t \in I} f_t(\mathbf{u}_t) + \sum_{t \in I} \|\mathbf{u}_t - \mathbf{u}_{t-1}\|_2$ when comparing with a sequence of time-varying comparators. Indeed, a similar phenomenon appears also in the generalization of small-loss static regret to dynamic regret (Zhao et al., 2024).

Lemma 10. Under Assumptions 1, 2, and 3, for any interval $[s_m, s_{m+1} - 1]$ determined by two consecutive intervals s_m and s_{m+1} , where we denote by s_m the m -th marker, Algorithm 3 ensures that for any comparators $\mathbf{u}_{s_m}, \dots, \mathbf{u}_{s_{m+1}-1} \in \mathcal{X}$, we have

$$2 \sum_{t=s_m}^{s_{m+1}-1} f_t(\mathbf{u}_t) + (126L + 5) \sum_{t=s_m+1}^{s_{m+1}-1} \|\mathbf{u}_t - \mathbf{u}_{t-1}\|_2 \geq \frac{1}{2}C_m,$$

where $C_m = \mathcal{G}(m)$ is the m -th threshold with the threshold function $\mathcal{G}(\cdot)$ defined at Eq. (19).

Lemma 11. Under Assumptions 1, 2, and 3, for any interval $I = [i, j] \in \tilde{\mathcal{C}}$ in the geometric covers and any time stamp $t \in [i, j]$, the variable m specified in Lemma 9 can be bounded by

$$m \leq \mathcal{O}\left(F_{[1,t]}^{\mathbf{u}} + P_{[1,t]}\right),$$

for any comparators $\mathbf{u}_1, \dots, \mathbf{u}_t \in \mathcal{X}$. This result can further imply that Algorithm 3 satisfies

$$\sum_{\tau=i}^t f_{\tau}(\mathbf{x}_{\tau}) - \sum_{\tau=i}^t f_{\tau}(\mathbf{u}_{\tau}) \leq \mathcal{O}\left(\sqrt{F_{[i,t]}^{\mathbf{u}}\left(P_{[i,t]} + \log\left(\min_{\mathbf{u}_{1:t}}\{F_{[1,t]}^{\mathbf{u}} + P_{[1,t]}\}\right)\right)} + P_{[i,t]}\right),$$

where $F_{[i,t]}^{\mathbf{u}} = \sum_{\tau=i}^t f_{\tau}(\mathbf{u}_{\tau})$ denotes the cumulative loss of the comparators, $P_{[i,t]} = \sum_{\tau=i+1}^t \|\mathbf{u}_{\tau} - \mathbf{u}_{\tau-1}\|_2$ is the path-length of comparators within the interval $[i, t]$, and $\min_{\mathbf{u}_{1:t}}\{F_I^{\mathbf{u}} + P_I\} = \min_{\mathbf{u}_1, \dots, \mathbf{u}_t} \left\{ \sum_{\tau=1}^t f_{\tau}(\mathbf{u}_{\tau}) + \sum_{\tau=2}^t \|\mathbf{u}_{\tau} - \mathbf{u}_{\tau-1}\|_2 \right\}$.

D.3 Proof of Theorem 5

Proof. The proof closely follows that of Theorem 4, with only the key steps outlined here. For a comprehensive understanding of the techniques, readers may check the proofs in Appendix C.2, where the problem setting is simpler with only a fixed comparator considered.

By a standard argument, we can identify markers s_p and s_q such that $s_{p-1} \leq r < s_p$ and $s_q \leq s < s_{q+1}$. Drawing upon Lemma 19, a sequence of intervals I_1, \dots, I_v exists where $i_1 = p$, $i_v \leq q < i_{v+1}$, and $v \leq \lceil \log_2(q - p + 2) \rceil$.

For the interval $[r, s_p - 1]$, we can directly bound the regret by the threshold of the cumulative loss and obtain that $\sum_{t=r}^{s_p-1} f_t(\mathbf{x}_t) - \sum_{t=r}^{s_p-1} f_t(\mathbf{u}_t) \leq \sum_{t=s_{p-1}}^{s_p-1} f_t(\mathbf{x}_t) \leq C_{p-1} + GD$, due to the threshold mechanism and the non-negative property of loss functions. The threshold generating function specified in Eq. (19) further indicates $C_{p-1} \leq \mathcal{O}(\log(F_T + P_T))$.

By Lemma 11, since I_1 to I_v belong to the covering intervals, then our algorithm can enjoy $\mathcal{O}(P_{I_k} + \sqrt{F_{I_k}(P_{I_k} + \log(P_{I_k} + F_{I_k})))$ regret bound for $k \in [v]$. Summing up the interval dynamic regret from I_1 to I_v gives

$$\begin{aligned} \sum_{t=s_p}^s f_t(\mathbf{x}_t) - \sum_{t=s_p}^s f_t(\mathbf{u}_t) &= \sum_{k=1}^{v-1} \sum_{t \in I_k} (f_t(\mathbf{x}_t) - f_t(\mathbf{u}_t)) + \sum_{t \in [s_{i_v}, s]} (f_t(\mathbf{x}_t) - f_t(\mathbf{u}_t)) \\ &\leq \sum_{k=1}^{v-1} \mathcal{O} \left(P_{I_k} + \sqrt{F_{I_k}^{\mathbf{u}} \left(P_{I_k} + \log \left(\min_{\mathbf{u}_{1:T}} \{F_T^{\mathbf{u}} + P_T\} \right) \right)} \right) \\ &\quad + \mathcal{O} \left(P_{[i_v, s]} + \sqrt{F_{[i_v, s]}^{\mathbf{u}} \left(P_{[i_v, s]} + \log \left(\min_{\mathbf{u}_{1:T}} \{F_T^{\mathbf{u}} + P_T\} \right) \right)} \right) \\ &\leq \mathcal{O} \left(P_I + \sqrt{F_I^{\mathbf{u}} \left(P_I + \log \left(\min_{\mathbf{u}_{1:T}} \{F_T^{\mathbf{u}} + P_T\} \right) \right)} \cdot v \right), \end{aligned}$$

where the last inequality makes use of Cauchy-Schwarz inequality and v denotes the number of combined intervals. With Lemma 10, we can upper bound v through following steps,

$$\begin{aligned} \min_{\mathbf{u}_{s_p:s_q-1}} \{F_{[s_p, s_q-1]}^{\mathbf{u}} + P_{[s_p, s_q-1]}\} &\geq \frac{1}{4} \sum_{i=p}^{q-1} C_i \geq \frac{C_1}{4} (q - p) \\ \Rightarrow v \leq \lceil \log_2(q - p + 2) \rceil &\leq \mathcal{O} \left(\log \left(\min_{\mathbf{u}_{r:s}} \{F_I^{\mathbf{u}} + P_I\} \right) \right). \end{aligned}$$

Combining the upper bounds of C_{p-1} and v yields the following interval dynamic regret,

$$\begin{aligned} \sum_{t=r}^s f_t(\mathbf{x}_t) - \sum_{t=q}^s f_t(\mathbf{u}_t) &\leq \mathcal{O} \left(P_I + \sqrt{F_I^{\mathbf{u}} \left(P_I + \log \left(\min_{\mathbf{u}_{1:T}} \{F_T^{\mathbf{u}} + P_T\} \right) \right)} \cdot v \right) + C_{p-1} + GD \\ &\leq \mathcal{O} \left(P_I + \sqrt{F_I^{\mathbf{u}} \left(P_I + \log \left(\min_{\mathbf{u}_{1:T}} \{F_T^{\mathbf{u}} + P_T\} \right) \right)} \cdot \log \left(\min_{\mathbf{u}_{r:s}} \{F_I^{\mathbf{u}} + P_I\} \right) \right) + \mathcal{O} \left(\log \left(\min_{\mathbf{u}_{1:T}} \{F_T^{\mathbf{u}} + P_T\} \right) \right) \\ &= \mathcal{O} \left(\sqrt{(F_I^{\mathbf{u}} + P_I) \left(P_I + \log \left(\min_{\mathbf{u}_{1:T}} \{F_T^{\mathbf{u}} + P_T\} \right) \right)} \cdot \log \left(\min_{\mathbf{u}_{r:s}} \{F_I^{\mathbf{u}} + P_I\} \right) \right). \end{aligned}$$

We mention that using a similar worst-case regret analysis to that in Appendix C.2 can ensure a safety guarantee, which can strictly match the worst-case interval dynamic regret bound in (Zhang et al., 2020). Details are omitted here. \square

D.4 Proof of Lemma 8

Proof. The proof is closely analogous to that of Theorem 3, except that this result can enjoy the anytime dynamic regret. For any time stamp $t \in [i, j]$, we can decompose the dynamic regret into meta-regret and base-regret, and bound them respectively. Notice that, we add the prefix “base:” to indicate that the regret analysis is over the base-algorithm level, as Algorithm 3 actually has three layers.

$$\sum_{\tau=i}^t \langle \nabla g_\tau(\mathbf{y}_\tau), \mathbf{y}_{\tau,m} - \mathbf{u}_\tau \rangle = \underbrace{\sum_{\tau=i}^t \langle \nabla g_\tau(\mathbf{y}_\tau), \mathbf{y}_{\tau,m} - \mathbf{y}_{\tau,m,k} \rangle}_{\text{base:meta-regret}} + \underbrace{\sum_{\tau=i}^t \langle \nabla g_\tau(\mathbf{y}_\tau), \mathbf{y}_{\tau,m,k} - \mathbf{u}_\tau \rangle}_{\text{base:base-regret}}, \quad (64)$$

where $\mathbf{y}_{\tau,m,k}$ denotes the k -th decision maintained by the m -th base-learner in Algorithm 3.

Upper bound of base:meta-regret. Notice that the meta-algorithm used for our dynamic algorithm is Hedge with self-confident tuning learning rates, so we have

$$\begin{aligned} \sum_{\tau=i}^t \langle \nabla g_\tau(\mathbf{y}_\tau), \mathbf{y}_{\tau,m} - \mathbf{y}_{\tau,m,k} \rangle &\leq 3D \sqrt{\ln N \sum_{\tau=i}^t \|\nabla g_\tau(\mathbf{y}_\tau)\|_2^2} + \frac{(6 + G^2 D^2) \sqrt{\ln N}}{2} \\ &\leq 6D \sqrt{L \ln N \sum_{\tau=i}^t f_\tau(\mathbf{x}_\tau)} + \frac{(6 + G^2 D^2) \sqrt{\ln N}}{2}. \end{aligned} \quad (65)$$

for any base:base-learner $k \in [N]$. The above reasoning is similar to Eq. (31).

Upper bound of base:base-regret. By Theorem 8 and Lemma 14, it is easy to verify once the learning rate is set, OGD ensures the following dynamic bound before tuning,

$$\begin{aligned} \sum_{\tau=i}^t \langle \nabla g_\tau(\mathbf{y}_\tau), \mathbf{y}_{\tau,m,k} - \mathbf{u}_\tau \rangle &\leq \frac{5D^2}{2\eta_k} + \frac{D}{\eta_k} \sum_{\tau=i+1}^t \|\mathbf{u}_{\tau-1} - \mathbf{u}_\tau\|_2 + \eta_k \sum_{\tau=i}^t \|\nabla g_\tau(\mathbf{y}_\tau)\|_2^2 \\ &\leq \frac{5D^2}{2\eta_k} + \frac{D}{\eta_k} \sum_{\tau=i+1}^t \|\mathbf{u}_{\tau-1} - \mathbf{u}_\tau\|_2 + 4\eta_k L \sum_{\tau=i}^t f_\tau(\mathbf{x}_\tau), \end{aligned} \quad (66)$$

for any base:base-learner $k \in [N]$.

Upper bound of anytime dynamic regret. Plugging (65) and (66) into (64), we can obtain the dynamic regret by tracking k -th base-learner,

$$\begin{aligned} &\sum_{\tau=i}^t \langle \nabla g_\tau(\mathbf{y}_\tau), \mathbf{y}_{\tau,m} - \mathbf{u}_\tau \rangle \\ &\leq 6D \sqrt{L \ln N \sum_{\tau=i}^t f_\tau(\mathbf{x}_\tau)} + \frac{5D^2 + 2DP_{[i,t]}}{2\eta_k} + 4\eta_k L \sum_{\tau=i}^t f_\tau(\mathbf{x}_\tau) + \frac{(6 + G^2 D^2) \sqrt{\ln N}}{2}. \end{aligned} \quad (67)$$

Next, we specific the base-learner to compare with. We are aiming to choose the one with a step size closest to the (near-)optimal step size till time t , $\eta_t^* = \sqrt{\frac{5D^2 + 2DP_t}{1 + 8LF_{[i,t]}^*}}$, where

we denote $F_{[i,t]}^{\mathbf{x}} = \sum_{\tau=i}^t f_{\tau}(\mathbf{x}_{\tau})$. With the same argument as the proof of Theorem 3, we can identify the base-learner k satisfying $\eta_k \leq \eta_t^* \leq 2\eta_k$. Tuning Eq. (67) with learning rate η_k specified above demonstrates that the dynamic regret can be upper bounded by

$$6D\sqrt{L \ln N \sum_{\tau=1}^t f_{\tau}(\mathbf{x}_{\tau})} + 3\sqrt{2(5D^2 + 2DP_{[i,t]})(1 + LF_{[i,t]}^{\mathbf{x}})} + \frac{(6 + G^2D^2)\sqrt{\ln N}}{2}. \quad (68)$$

This ends the proof. \square

D.5 Proof of Lemma 9

Proof. By Theorem 1 and the combination of our algorithm, we can upper bound the interval dynamic on interval I into two terms as before,

$$\begin{aligned} \sum_{\tau=i}^t f_{\tau}(\mathbf{x}_{\tau}) - \sum_{\tau=i}^t f_{\tau}(\mathbf{u}_{\tau}) &\leq \sum_{\tau=i}^t \langle \nabla g_{\tau}(\mathbf{y}_{\tau}), \mathbf{y}_{\tau} - \mathbf{u}_{\tau} \rangle \\ &= \sum_{\tau=i}^t \underbrace{\langle \nabla g_{\tau}(\mathbf{y}_{\tau}), \mathbf{y}_{\tau} - \mathbf{y}_{\tau,m} \rangle}_{\text{meta-regret}} + \underbrace{\langle \nabla g_{\tau}(\mathbf{y}_{\tau}), \mathbf{y}_{\tau,m} - \mathbf{u}_{\tau} \rangle}_{\text{base-regret}}, \end{aligned}$$

where the base-learner's decision $\mathbf{y}_{\tau,m}$ comes from m -th base-learner, namely the efficient dynamic algorithm, which is also a combination of several OGD algorithms and active on the considered interval. Our algorithm is a three-layer structure indeed, but we hide the details of the efficient dynamic algorithm by Lemma 8.

Upper bound of meta-regret. Our interval dynamic algorithm uses essentially the same meta-algorithm and cover as efficient adaptive algorithm (see Algorithm 2), so we can directly use Lemma 4 to upper bound the meta-regret,

$$\sum_{\tau=i}^t \langle \nabla g_{\tau}(\mathbf{y}_{\tau}), \mathbf{y}_{\tau} - \mathbf{y}_{\tau,m} \rangle \leq 4D \left(3\sqrt{\ln(1+2m)} + \mu_t \right) \sqrt{LF_{[i,t]}^{\mathbf{x}}} + 18GD \ln(1+2m) + 6GD\mu_t,$$

where we denote $\mu_t = \ln(1 + (1 + \ln(1+t))/(2e))$ and $F_{[a,b]}^{\mathbf{x}} = \sum_{\tau=a}^b f_{\tau}(\mathbf{x}_{\tau})$.

Upper bound of base-regret. By the step size pool setting and Lemma 8, we know that for any $t \in [i, j]$, our base-algorithm ensures anytime dynamic regret,

$$\begin{aligned} \sum_{\tau=i}^t \langle \nabla g_{\tau}(\mathbf{y}_{\tau}), \mathbf{y}_{\tau,m} - \mathbf{u}_{\tau} \rangle &\leq 6D\sqrt{L \ln N F_{[i,t]}^{\mathbf{x}}} + 3\sqrt{2L(5D^2 + 2DP_t)F_{[i,t]}^{\mathbf{x}}} \\ &\quad + 3\sqrt{2(5D^2 + 2DP_{[i,t]})} + \frac{(6 + G^2D^2)\sqrt{\ln N}}{2}. \end{aligned}$$

Upper bound of interval dynamic regret. Combining the meta-regret and base-regret discussed above, applying Lemma 25 and omitting tedious calculations, we have

$$\sum_{\tau=i}^t f_{\tau}(\mathbf{x}_{\tau}) - f_{\tau}(\mathbf{u}_{\tau})$$

$$\begin{aligned}
 &\leq \left(12D\sqrt{\ln(1+2m)} + 4D\mu_t + 6D\sqrt{\ln N} + 3\sqrt{2(5D^2 + 2DP_{[i,t]})}\right) \sqrt{LF_{[i,t]}^{\mathbf{u}}} \\
 &\quad + 27GD\ln(1+2m) + 9GD\mu_t + \frac{9}{2}\sqrt{2(5D^2 + 2DP_{[i,t]})} + \frac{3(6 + G^2D^2)\sqrt{\ln N}}{4} \\
 &\quad + \frac{3L}{2} \left(12D\sqrt{\ln(1+2m)} + 4D\mu_t + 6D\sqrt{\ln N} + 3\sqrt{2(5D^2 + 2DP_{[i,t]})}\right)^2 \\
 &\leq \mathcal{O}\left(\sqrt{F_{[i,t]}^{\mathbf{u}}(P_{[i,t]} + \log m)} + P_{[i,t]} + \log m\right).
 \end{aligned}$$

□

D.6 Proof of Lemma 10

Proof. We introduce the notation $F_{[a,b]} = \sum_{t=a}^b f_t(\mathbf{u}_t)$ to denote the cumulative loss of comparators and $P_{[a,b]} = \sum_{t=a+1}^b \|\mathbf{u}_t - \mathbf{u}_{t-1}\|_2$ to denote the path length of comparators. By Lemma 9 and derivations, we can isolate F_I and P_I from the square root in (63), where we choose $I = [s_m, s_{m+1} - 1]$ as there always exists an interval in the schedule to cover it:

$$\sum_{t \in I} f_t(\mathbf{x}_t) \leq 2 \sum_{t \in I} f_t(\mathbf{u}_t) + (126L + 5)DP_I + \frac{1}{2}\mathcal{G}(m),$$

where $\mathcal{G}(\cdot)$ is the threshold function defined at Eq. (19). By the threshold mechanism, we know that the cumulative loss within I exceeds C_m , which implies

$$2F_I^{\mathbf{u}} + (126L + 5)DP_I \geq \sum_{t \in I} f_t(\mathbf{x}_t) - \frac{1}{2}\mathcal{G}(m) \geq C_m - \frac{1}{2}\mathcal{G}(m) = \frac{1}{2}C_m,$$

where we denote by $F_I^{\mathbf{u}} = \sum_{t \in I} f_t(\mathbf{u}_t)$.

□

D.7 Proof of Lemma 11

Proof. We assume that the m -th base-learner is initialized at the beginning of interval $[i, j] \in \tilde{\mathcal{C}}$ (in other words, $i = s_m$), where s_m denotes the m -th marker. Before time stamp i , the schedule has registered $m - 1$ markers, i.e., from s_1 to s_{m-1} . By Lemma 10, for any comparators, we have

$$\sum_{k=1}^{m-1} \left(2F_{[s_k, s_{k+1}-1]}^{\mathbf{u}} + (126L + 5)DP_{[s_k, s_{k+1}-1]}\right) \geq \frac{1}{2} \sum_{u=1}^{m-1} C_u \geq \frac{C_1}{2}(m-1).$$

Rearranging the above inequality provides the upper bound of m as

$$\begin{aligned}
 m &\leq 1 + \frac{2}{C_1} \left(\sum_{u=1}^{m-1} 2F_{[s_u, s_{u+1}-1]}^{\mathbf{u}} + (126L + 5)DP_{[s_u, s_{u+1}-1]} \right) \\
 &\leq 1 + \frac{2}{C_1} \left(2F_{[s_1, s_m-1]}^{\mathbf{u}} + (126L + 5)DP_{[s_1, s_m-1]} \right) \\
 &\leq 1 + \frac{2}{C_1} \left(2F_{[1, t]}^{\mathbf{u}} + (126L + 5)DP_{[1, t]} \right).
 \end{aligned}$$

Substituting the upper bound of m into Lemma 9 finishes the proof.

□

Appendix E. Omitted Details for Efficient Projection Examples

In this section, we present the omitted details for the two applications of our proposed efficient projection scheme.

E.1 Online Non-stochastic Control

In this subsection, we provide the required assumptions for online non-stochastic control and then provide the proof of Theorem 6.

E.1.1 ASSUMPTIONS

The following assumptions are required by Theorem 6, which are commonly used in the non-stochastic control analysis (Hazan and Singh, 2022; Zhao et al., 2023).

Assumption 4. The system matrices are bounded, i.e., $\|A\|_{\text{op}} \leq \kappa_A$ and $\|B\|_{\text{op}} \leq \kappa_B$. Besides, the disturbance $\|w_t\| \leq W$ holds for any $t \in [T]$.

Assumption 5. The cost function $c_t(x, u)$ is convex. Further, when $\|x\|, \|u\| \leq D$, it holds that $|c_t(x, u)| \leq \beta D^2$ and $\|\nabla_x c_t(x, u)\|, \|\nabla_u c_t(x, u)\| \leq G_c D$.

Assumption 6. DAC controller $\pi(K, M)$ satisfies

1. K is (κ, γ) -strongly stable, i.e., there exist matrices L, H satisfying $A - BK = HLH^{-1}$, such that,
 - (a) The spectral norm of L satisfies $\|L\| \leq 1 - \gamma$.
 - (b) The controller and transforming matrices are bounded, i.e., $\|K\| \leq \kappa$ and $\|H\|, \|H^{-1}\| \leq \kappa$.
2. $M \in \mathcal{M}$ such that $\left\{ M = (M^{[1]}, \dots, M^{[H]}) \in (\mathbb{R}^{d_u \times d_x})^H \mid \|M^{[i]}\|_{\text{op}} \leq \kappa_B \kappa^3 (1 - \gamma)^i \right\}$.

E.1.2 PROOF OF THEOREM 6

The challenge in proving Theorem 6 lies in accounting for the switching-cost while improving the efficiency. The crucial observation is given by $\|M_{t-1} - M_t\|_{\text{F}} \leq \|M'_{t-1} - M'_t\|_{\text{F}}$. This relationship is derived by the nonexpanding property of projection operator (Nemirovski et al., 2009). This implies that the switching-cost within the original domain can be constrained by that in the surrogate domain, which the algorithm is designed to minimize.

Proof. The proof mainly follows the one of Scream.Control. We present the essential steps to demonstrate the application of efficient reduction here and refer interested readers to Appendix C.2.3 in Zhao et al. (2022b) for comprehensive proof.

We denote by $f_t(\cdot) : \mathcal{M}^{H+2} \mapsto \mathbb{R}$ the truncated loss and the dynamic regret enjoys the following decomposition:

$$\sum_{t=1}^T c_t(x_t, u_t) - \sum_{t=1}^T c_t(x_t^{\pi_t}, u_t^{\pi_t})$$

$$\begin{aligned}
 &= \sum_{t=1}^T c_t(x_t^K(M_{0:t-1}), u_t^K(M_{0:t})) - \sum_{t=1}^T c_t(x_t^K(M_{0:t-1}^*), u_t^K(M_{0:t}^*)) \\
 &= \underbrace{\sum_{t=1}^T c_t(x_t^K(M_{0:t-1}), u_t^K(M_{0:t})) - \sum_{t=1}^T f_t(M_{t-1-H:t})}_{A_T} + \underbrace{\sum_{t=1}^T f_t(M_{t-1-H:t}) - \sum_{t=1}^T f_t(M_{t-1-H:t}^*)}_{B_T} \\
 &\quad + \underbrace{\sum_{t=1}^T f_t(M_{t-1-H:t}^*) - \sum_{t=1}^T c_t(x_t^K(M_{0:t-1}^*), u_t^K(M_{0:t}^*))}_{C_T}.
 \end{aligned}$$

Notice that A_T and C_T represent the approximation error induced by the truncated loss, which does not involve the efficient reduction and can be bounded effectively. As for B_T :

$$\begin{aligned}
 B_T &\leq \sum_{t=1}^T \tilde{f}_t(M_t) - \sum_{t=1}^T \tilde{f}_t(M_t^*) + \lambda \sum_{t=2}^T \|M_{t-1} - M_t\|_F + \lambda \sum_{t=2}^T \|M_{t-1}^* - M_t^*\|_F \\
 &\leq \sum_{t=1}^T \langle \nabla \tilde{f}_t(M_t), M_t - M_t^* \rangle + \lambda \sum_{t=2}^T \|M_{t-1} - M_t\|_F + \lambda \sum_{t=2}^T \|M_{t-1}^* - M_t^*\|_F \\
 &\leq \sum_{t=1}^T \langle \nabla g_t(M_t'), M_t' - M_t^* \rangle + \lambda \sum_{t=2}^T \|M_{t-1} - M_t\|_F + \lambda \sum_{t=2}^T \|M_{t-1}^* - M_t^*\|_F \\
 &\leq \sum_{t=1}^T \langle \nabla g_t(M_t'), M_t' - M_t^* \rangle + \lambda \sum_{t=2}^T \|M_{t-1}' - M_t'\|_F + \lambda \sum_{t=2}^T \|M_{t-1}^* - M_t^*\|_F,
 \end{aligned}$$

where $\lambda = (H+2)^2 L_f$ is a constant. The first inequality is by the coordinate-Lipschitz continuity of truncated function $f_t(\cdot)$. The third inequality is by the reduction mechanism and the final inequality is by $\|M_{t-1} - M_t\|_F \leq \|M_{t-1}' - M_t'\|_F$ derived from the nonexpanding property of projection (Nemirovski et al., 2009) and that one can verify in general this property holds for nearest point projection in Hilbert space.

Remind that Scream.Control employed in Algorithm 4 aims at minimizing the dynamic regret with switching-cost in domain \mathcal{M}' , which can guarantee $\tilde{\mathcal{O}}(\sqrt{T(1+P_T)})$ regret bound by Theorem 4 in Zhao et al. (2022b). Thus, by taking into account that $\|\nabla g_t(M_t')\|_F \leq \|\nabla \tilde{f}_t(M_t)\|_F$ by the efficient reduction mechanism, which is true under non-stochastic control setting since the algorithm optimizes the linearized loss and employs the Frobenius norm as the projection distance metric, we can derive our result. \square

E.2 Online Principal Component Analysis

This section provides omitted details for the online PCA problem. In Appendix E.2.1 we provide the guarantee for base-algorithm and the lemma justifying the projection complexity. Appendix E.2.2 presents the overall proof of Theorem 7.

E.2.1 KEY LEMMAS

The following lemma presents the base-regret for the employed gradient-based algorithm.

Lemma 12. *Assuming $\|\mathbf{X}_t\|_F \leq 1$ and $k \leq \frac{d}{2}$, then any base-algorithm employed in Algorithm 5 specified as,*

$$\hat{\mathbf{P}}_{t+1,i}^{s'} = \hat{\mathbf{P}}_{t,i}^s - \eta_i \nabla g_t(\hat{\mathbf{P}}_t^s), \quad \hat{\mathbf{P}}_{t+1,i}^s = \hat{\mathbf{P}}_{t+1,i}^{s'} \left(\mathbb{1}_{\{\|\hat{\mathbf{P}}_{t+1,i}^{s'}\|_F \leq \sqrt{k}\}} + \frac{\sqrt{k}}{\|\hat{\mathbf{P}}_{t+1,i}^{s'}\|_F} \mathbb{1}_{\{\|\hat{\mathbf{P}}_{t+1,i}^{s'}\|_F > \sqrt{k}\}} \right),$$

which is active during time span $I = [r, s] \subseteq [T]$ and is indexed by number $i \in [T]$, ensures the following regret bound for any comparator $\mathbf{P} \in \mathcal{P}_k$ by tuning learning rate as $\eta = \frac{k(d-k)}{d|I|}$,

$$\sum_{t=r}^s \text{tr} \left(\nabla g_t(\hat{\mathbf{P}}_t^s) \cdot \hat{\mathbf{P}}_{t,i}^s \right) - \sum_{t=r}^s \text{tr} \left(\nabla g_t(\hat{\mathbf{P}}_t^s) \cdot \mathbf{P} \right) \leq \mathcal{O} \left(\sqrt{k \cdot |I|} \right),$$

The above claim can be verified by the proof in Appendix. B of Nie et al. (2016) together with $\|\nabla g_t(\hat{\mathbf{P}}_t^s)\|_F \leq \|\nabla f_t(\hat{\mathbf{P}}_t^s)\|_F$, which can be verified by noticing that the loss function $f_t(\mathbf{P})$ is coordinate-wise linear with \mathbf{P} and we use Frobenius norm as the distance metric.

The following lemma provides the details to project decision onto domain $\hat{\mathcal{P}}_k$.

Lemma 13 (Lemma 3.2 of Arora et al. (2013)). *Let $\mathbf{P}' \in \mathbb{R}^{d \times d}$ be a symmetric matrix, with eigenvalues $\sigma'_1, \dots, \sigma'_d$ and associated eigenvectors $\mathbf{v}'_1, \dots, \mathbf{v}'_d$. Its projection $\mathbf{P} = \Pi_{\hat{\mathcal{P}}_k}[\mathbf{P}']$ onto the domain $\hat{\mathcal{P}}_k$ with respect to the Frobenius norm, is the unique feasible matrix which has the same eigenvectors as \mathbf{P}' , with the associated eigenvalues $\sigma_1, \dots, \sigma_d$ satisfying:*

$$\sigma_i = \max(0, \min(1, \sigma'_i + S)), i \in [d],$$

with $S \in \mathbb{R}$ being chosen in such a way that $\sum_{i=1}^d \sigma_i = k$. Moreover, there exists an algorithm to find S in an $\mathcal{O}(d \log d)$ running time complexity.

E.2.2 PROOF OF THEOREM 7

The proof of Theorem 7 enjoys much similarity as the one of efficient adaptive algorithm. We refer the readers to Appendix 3 for more details.

Proof. We mainly present the key steps for applying our reduction. For any interval $I = [r, s] \subseteq [T]$ and any comparator $\mathbf{P} \in \mathcal{P}_k$, starting with the linearity of expectation, we have,

$$\begin{aligned} \sum_{t=r}^s \mathbb{E}[f_t(\mathbf{P}_t)] - f_t(\mathbf{P}) &= \sum_{t=r}^s \mathbb{E}[\text{tr}((\mathbf{I} - \mathbf{P}_t)\mathbf{X}_t)] - \text{tr}((\mathbf{I} - \mathbf{P})\mathbf{X}_t) \\ &= \sum_{t=r}^s \text{tr}((\mathbf{I} - \hat{\mathbf{P}}_t)\mathbf{X}_t) - \text{tr}((\mathbf{I} - \mathbf{P})\mathbf{X}_t) \\ &= \sum_{t=r}^s \text{tr}(\nabla f_t(\hat{\mathbf{P}}_t) \cdot \hat{\mathbf{P}}_t) - \text{tr}(\nabla f_t(\hat{\mathbf{P}}_t) \cdot \mathbf{P}) \\ &\leq \sum_{t=r}^s \text{tr}(\nabla g_t(\hat{\mathbf{P}}_t^s) \cdot \hat{\mathbf{P}}_t^s) - \text{tr}(\nabla g_t(\hat{\mathbf{P}}_t^s) \cdot \mathbf{P}) \end{aligned}$$

$$\begin{aligned}
 &= \underbrace{\sum_{t=r}^s \text{tr} \left(\nabla g_t(\hat{\mathbf{P}}_t^s) \cdot \hat{\mathbf{P}}_t^s \right) - \sum_{t=r}^s \text{tr} \left(\nabla g_t(\hat{\mathbf{P}}_t^s) \cdot \hat{\mathbf{P}}_{t,i}^s \right)}_{\text{meta-regret}} \\
 &\quad + \underbrace{\sum_{t=r}^s \text{tr} \left(\nabla g_t(\hat{\mathbf{P}}_t^s) \cdot \hat{\mathbf{P}}_{t,i}^s \right) - \sum_{t=r}^s \text{tr} \left(\nabla g_t(\hat{\mathbf{P}}_t^s) \cdot \mathbf{P} \right)}_{\text{base-regret}},
 \end{aligned}$$

where the first inequality is by Theorem 1, which is true under online PCA setting, since the optimization operates within the Hilbert space.

Since we employ Adapt-ML-Prod and standard geometric covering schedule to ensemble the base-learners, then one can expect that $\text{meta-regret} \leq \mathcal{O}(\sqrt{k \cdot |I| \cdot \log T})$. By Lemma 12, the base-regret is of order $\text{base-regret} \leq \mathcal{O}(\sqrt{k \cdot |I|})$. Combining these two bounds together, we finish the proof. \square

Appendix F. Useful Lemmas

This section collects some lemmas useful for the proofs.

F.1 OGD and Dynamic Regret

This part provides the dynamic regret of online gradient descent (OGD) (Zinkevich, 2003) and scale-free online gradient descent (SOGD) (Orabona and Pál, 2018) from the view of online mirror descent (OMD), which is a common and powerful online learning framework. Following the analysis in (Zhao et al., 2024), we can obtain dynamic regret of OGD and SOGD in a unified view owing to the versatility of OMD. Specifically, OMD updates by

$$\mathbf{x}_{t+1} = \arg \min_{\mathbf{x} \in \mathcal{X}} \eta_t \langle \nabla f_t(\mathbf{x}_t), \mathbf{x} \rangle + \mathcal{D}_\psi(\mathbf{x}, \mathbf{x}_t), \quad (69)$$

where $\eta_t > 0$ is the time-varying step size, $f_t(\cdot) : \mathbf{x} \mapsto \mathbb{R}$ is the convex loss function, and $\mathcal{D}_\psi(\cdot, \cdot)$ is the Bregman divergence induced by $\psi(\cdot)$ defined as $\mathcal{D}_\psi(\mathbf{x}, \mathbf{y}) = \psi(\mathbf{x}) - \psi(\mathbf{y}) - \langle \nabla \psi(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle$. OMD enjoys the following dynamic regret guarantee (Zhao et al., 2024).

Theorem 8 (Theorem 1 of Zhao et al. (2024)). *Suppose that the regularizer $\psi : \mathcal{X} \mapsto \mathbb{R}$ is 1-strongly convex with respect to the norm $\|\cdot\|$. The dynamic regret of Optimistic Mirror Descent (OMD) whose update rule specified in (69) is bounded as follows:*

$$\begin{aligned}
 &\sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}_t) \\
 &\leq \sum_{t=1}^T \eta_t \|\nabla f_t(\mathbf{x}_t)\|_*^2 + \sum_{t=1}^T \frac{1}{\eta_t} \left(\mathcal{D}_\psi(\mathbf{u}_t, \mathbf{x}_t) - \mathcal{D}_\psi(\mathbf{u}_t, \mathbf{x}_{t+1}) \right) - \sum_{t=1}^T \frac{1}{\eta_t} \mathcal{D}_\psi(\mathbf{x}_{t+1}, \mathbf{x}_t),
 \end{aligned}$$

which holds for any comparator sequence $\mathbf{u}_1, \dots, \mathbf{u}_T \in \mathcal{X}$.

Choosing $\psi(\mathbf{x}) = \frac{1}{2}\|\mathbf{x}\|_2^2$ will lead to the update form of online gradient descent used as base-learners in our algorithm:

$$\mathbf{x}_{t+1} = \arg \min_{\mathbf{x} \in \mathcal{X}} \eta_t \langle \nabla f_t(\mathbf{x}_t), \mathbf{x} \rangle + \frac{1}{2} \|\mathbf{x} - \mathbf{x}_t\|_2^2, \quad (70)$$

where the Bregman divergence becomes $\mathcal{D}_\psi(\mathbf{x}, \mathbf{x}_t) = \frac{1}{2} \|\mathbf{x} - \mathbf{x}_t\|_2^2$. We proceed to show the dynamic regret of online gradient descent (OGD),

Lemma 14. *Under Assumption 2, by choosing static step size $\eta_t = \eta > 0$, Online Gradient Descent defined in Eq. (70) satisfies:*

$$\sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}_t) \leq \frac{7D^2}{4\eta} + \frac{D}{\eta} \sum_{t=2}^T \|\mathbf{u}_{t-1} - \mathbf{u}_t\|_2 + \eta \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t)\|_2^2$$

for any comparator sequence $\mathbf{u}_1, \dots, \mathbf{u}_T \in \mathcal{X}$.

Proof. Applying Theorem 8 with the $\psi(\mathbf{x}) = \frac{1}{2}\|\mathbf{x}\|_2^2$ and fixed step size $\eta_t = \eta > 0$ gives

$$\begin{aligned} & \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}_t) \\ & \leq \frac{1}{2\eta} \sum_{t=1}^T \left(\|\mathbf{u}_t - \mathbf{x}_t\|_2^2 - \|\mathbf{u}_t - \mathbf{x}_{t+1}\|_2^2 \right) + \eta \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t)\|_2^2 - \frac{1}{2\eta} \sum_{t=1}^T \|\mathbf{x}_t - \mathbf{x}_{t+1}\|_2^2 \\ & \leq \frac{1}{2\eta} \sum_{t=1}^T \left(\|\mathbf{x}_t\|_2^2 - \|\mathbf{x}_{t+1}\|_2^2 \right) + \frac{1}{\eta} \sum_{t=1}^T (\mathbf{x}_{t+1} - \mathbf{x}_t)^\top \mathbf{u}_t + \eta \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t)\|_2^2 \\ & \leq \frac{1}{2\eta} \|\mathbf{x}_1\|_2^2 + \frac{1}{\eta} \left(\mathbf{x}_{T+1}^\top \mathbf{u}_T - \mathbf{x}_1^\top \mathbf{u}_1 \right) + \frac{1}{\eta} \sum_{t=2}^T (\mathbf{u}_{t-1} - \mathbf{u}_t)^\top \mathbf{x}_t + \eta \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t)\|_2^2 \\ & \leq \frac{7D^2}{4\eta} + \frac{D}{\eta} \sum_{t=2}^T \|\mathbf{u}_{t-1} - \mathbf{u}_t\|_2 + \eta \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t)\|_2^2, \end{aligned}$$

where the last inequality is due to the domain boundedness. This ends the proof. \square

Lemma 15 (Stability Lemma). *Suppose the regularizer $\psi : \mathcal{X} \mapsto \mathbb{R}$ is 1-strongly convex with respect to the norm $\|\cdot\|$. The subsequent decisions $\mathbf{x}_{t+1}, \mathbf{x}_t$ specialized in the OMD update rule (69) satisfy $\|\mathbf{x}_{t+1} - \mathbf{x}_t\| \leq \eta_t \|\nabla f_t(\mathbf{x}_t)\|_*$.*

F.2 Self-confident Tuning

Orabona and Pál (2018) proved the regret bound of SOGD, and for completeness, we here provide the regret analysis under the OMD framework. Indeed, SOGD can be treated as OMD with a self-confident learning rate. Thus, we have the following lemma.

Lemma 16. *Under Assumptions 1 and 2, the OMD algorithm defined in equation (69) with the choices of regularizer $\psi(\mathbf{x}) = \frac{1}{2}\|\mathbf{x}\|_2^2$ and time-varying learning rates $\eta_t = \frac{D}{2\sqrt{\delta + \sum_{s=1}^{t-1} \|\nabla f_s(\mathbf{x}_s)\|_2^2}}$ with $\delta > 0$ enjoys the following guarantee:*

$$\sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}) \leq 2D \cdot \sqrt{\delta + \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t)\|_2^2}.$$

Proof. Applying Theorem 8, we have that for any comparator $\mathbf{u} \in \mathcal{X}$,

$$\begin{aligned} & \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}_t) \\ & \leq \sum_{t=1}^T \frac{1}{2\eta_t} \left(\|\mathbf{u} - \mathbf{x}_t\|_2^2 - \|\mathbf{u} - \mathbf{x}_{t+1}\|_2^2 \right) + \sum_{t=1}^T \eta_t \|\nabla f_t(\mathbf{x}_t)\|_2^2 - \sum_{t=1}^T \frac{1}{2\eta_t} \|\mathbf{x}_t - \mathbf{x}_{t+1}\|_2^2 \\ & \leq \frac{1}{2\eta_1} \|\mathbf{u} - \mathbf{x}_1\|_2^2 + \sum_{t=2}^T \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) \frac{\|\mathbf{u} - \mathbf{x}_t\|_2^2}{2} + \sum_{t=1}^T \eta_t \|\nabla f_t(\mathbf{x}_t)\|_2^2 \\ & \leq \frac{D^2}{2\eta_1} + \frac{D^2}{2} \sum_{t=2}^T \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) + \sum_{t=1}^T \eta_t \|\nabla f_t(\mathbf{x}_t)\|_2^2 \\ & = \frac{D^2}{2\eta_T} + \sum_{t=1}^T \eta_t \|\nabla f_t(\mathbf{x}_t)\|_2^2. \end{aligned}$$

Then, applying Lemma 21 to the second term and using the definition of η_T , we obtain the following regret bound:

$$\begin{aligned} \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}_t) & \leq D \cdot \sqrt{\delta + \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t)\|_2^2} + D \left(\sqrt{\delta + \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t)\|_2^2} - \sqrt{\delta} \right) \\ & \leq 2D \cdot \sqrt{\delta + \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t)\|_2^2}, \end{aligned}$$

which completes the proof. \square

To bound the meta-regret of our dynamic methods, we introduce the FTRL lemma (Orabona, 2019, Corollary 7.8) under the time-varying learning rates.

Lemma 17 (FTRL Lemma). *Suppose that the regularizer function $\psi : \mathcal{X} \mapsto \mathbb{R}$ is α -strongly convex with respect to the norm $\|\cdot\|$. Let f_t be a sequence of convex loss functions and $\psi_t(\mathbf{x}) = \frac{1}{\eta_t}(\psi(\mathbf{x}) - \min_{\mathbf{x}' \in \mathcal{X}} \psi(\mathbf{x}'))$, where $\eta_{t+1} \leq \eta_t$ holds for $t \in [T]$. Then the decision sequence \mathbf{x}_t generated by the following FTRL update rule*

$$\mathbf{x}_t = \arg \min_{\mathbf{x} \in \mathcal{X}} \left\{ \psi_t(\mathbf{x}) + \sum_{\tau=1}^{t-1} f_\tau(\mathbf{x}) \right\}$$

satisfies the following regret upper bound for any $\mathbf{u} \in \mathcal{X}$,

$$\sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}) \leq \frac{\psi(\mathbf{u}) - \min_{\mathbf{x} \in \mathcal{X}} \psi(\mathbf{x})}{\eta_{T+1}} + \frac{1}{2\alpha} \sum_{t=1}^T \eta_t \|\nabla f_t(\mathbf{x}_t)\|_*^2.$$

Based on it, we can derive the regret upper bound for the Hedge algorithm with self-confident learning rates.

Lemma 18. *Consider the prediction with expert advice setting with N experts and the linear loss $f_t(\mathbf{x}) = \langle \ell_t, \mathbf{x} \rangle$, where $\ell_t \in \mathbb{R}^d$. Then the self-confident tuning Hedge, whose initial decision is $\mathbf{p}_1 = 1/N \cdot \mathbf{1}$ and update rules are*

$$p_{t+1,i} \propto \exp \left(\varepsilon_{t+1} \sum_{\tau=1}^t \ell_{\tau,i} \right) \text{ with } \varepsilon_{t+1} = \sqrt{\frac{\ln N}{1 + \sum_{\tau=1}^t \|\ell_{\tau}\|_{\infty}^2}}$$

ensures the following regret guarantee: for any $i \in [N]$

$$\sum_{t=1}^T \langle \mathbf{p}_t, \ell_t \rangle - \sum_{t=1}^T \ell_{t,i} \leq 3 \sqrt{\ln N \cdot \left(1 + \sum_{t=1}^T \|\ell_t\|_{\infty}^2 \right)} + \frac{\sqrt{\ln N}}{2} \cdot \max_{t \in [T]} \|\ell_t\|_{\infty}^2.$$

Proof. It is easy to verify that this Hedge update can be treated as a special case of the time-varying FTRL algorithm by choosing $\psi(\mathbf{p}) = \sum_{s=1}^N p_s \ln p_s$, which is 1-strongly convex with respect to $\|\cdot\|_1$, and $\psi_t(\mathbf{p}) = \frac{1}{\varepsilon_t} \psi(\mathbf{p})$. Thus, by Lemma 17, we have

$$\begin{aligned} \sum_{t=1}^T \langle \mathbf{p}_t, \ell_t \rangle - \sum_{t=1}^T \ell_{t,i} &\leq \frac{\ln N}{\varepsilon_{T+1}} + \frac{1}{2} \sum_{t=1}^T \varepsilon_t \|\ell_t\|_{\infty}^2 \\ &\leq \frac{\ln N}{\varepsilon_{T+1}} + \frac{\sqrt{\ln N}}{2} \cdot \left(4 \sqrt{1 + \sum_{t=1}^T \|\ell_t\|_{\infty}^2} + \max_{t \in [T]} \|\ell_t\|_{\infty}^2 \right) \\ &= 3 \sqrt{\ln N \cdot \left(1 + \sum_{t=1}^T \|\ell_t\|_{\infty}^2 \right)} + \frac{\sqrt{\ln N}}{2} \cdot \max_{t \in [T]} \|\ell_t\|_{\infty}^2, \end{aligned}$$

where the first inequality chooses \mathbf{u} as the one-hot vector with all entries being 0 except the i -th one as 1, and second inequality is by Lemma 23. \square

F.3 Facts on Geometric Covers

Lemma 19 (Lemma 11 of Zhang et al. (2019)). *Let $[s_p, s_q] \subseteq [T]$ be an arbitrary interval that starts from a marker s_p and ends at another marker s_q . Then, we can find a sequence of consecutive intervals $I_1 = [s_{i_1}, s_{i_2} - 1]$, $I_2 = [s_{i_2}, s_{i_3} - 1]$, \dots , $I_v = [s_{i_v}, s_{i_{v+1}} - 1] \in \tilde{\mathcal{C}}$ such that $i_1 = p$, $i_v \leq q < i_{v+1}$, and $v \leq \lceil \log_2(q - p + 2) \rceil$.*

F.4 Technical Lemmas

In this part, we present several technical lemmas used in the proofs.

Lemma 20 (Lemma 3.1 of [Srebro et al. \(2010\)](#)). *Suppose $f : \mathbb{R}^d \mapsto \mathbb{R}$ is an L -smooth function, then for any $\mathbf{x} \in \mathbb{R}^d$, we have*

$$\|\nabla f(\mathbf{x})\|_2^2 \leq 2L \cdot \left(f(\mathbf{x}) - \min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x}) \right). \quad (71)$$

Furthermore, when the function is non-negative, we have $\|\nabla f(\mathbf{x})\|_2 \leq \sqrt{2L \cdot f(\mathbf{x})}$.

Lemma 21 (Lemma 3.5 of [Auer et al. \(2002\)](#)). *Let l_1, \dots, l_T be non-negative real numbers. Then*

$$\sum_{t=1}^T \frac{l_t}{\sqrt{\delta + \sum_{i=1}^t l_i}} \leq 2 \left(\sqrt{\delta + \sum_{t=1}^T l_t} - \sqrt{\delta} \right).$$

Lemma 22 (Lemma 14 of [Gaillard et al. \(2014\)](#)). *Let $a_0 > 0$ and $a_1, \dots, a_m \in [0, 1]$ be real numbers and let $f : (0, +\infty) \mapsto [0, +\infty)$ be a non-increasing function. Then*

$$\sum_{i=1}^m a_i f(a_0 + \dots + a_{i-1}) \leq f(a_0) + \int_{a_0}^{a_0 + a_1 + \dots + a_m} f(u) \, du.$$

Lemma 23 (Lemma 4.8 of [Pogodin and Lattimore \(2019\)](#)). *Let a_1, a_2, \dots, a_T be non-negative real numbers. Then*

$$\sum_{t=1}^T \frac{a_t}{\sqrt{1 + \sum_{s=1}^{t-1} a_s}} \leq 4 \sqrt{1 + \sum_{t=1}^T a_t + \max_{t \in [T]} a_t}.$$

Lemma 24 (Lemma 5 of [Shalev-Shwartz \(2007\)](#)). *For any $x, y, a \in \mathbb{R}_+$ satisfying $x - y \leq \sqrt{ax}$, we have $x - y \leq \sqrt{ay} + a$.*

Based on Lemma 24, we can achieve the following variant.

Lemma 25. *For any $x, y, a, b \in \mathbb{R}_+$ satisfying $x - y \leq \sqrt{ax} + b$, $x - y \leq \sqrt{ay + ab} + a + b$.*

Lemma 26 (Lemma 13 of [Gaillard et al. \(2014\)](#)). *For all $x > 0$ and $\alpha \geq 1$, $x \leq x^\alpha + \frac{\alpha-1}{e}$.*